Vision HDL Toolbox™

User's Guide

MATLAB®

R2023a

MathWorks®

How to Contact MathWorks

Latest news: www.mathworks.com

Sales and services: www.mathworks.com/sales_and_services

User community: www.mathworks.com/matlabcentral

Technical support: www.mathworks.com/support/contact_us

Phone: 508-647-7000

The MathWorks, Inc.
1 Apple Hill Drive
Natick, MA 01760-2098

Vision HDL Toolbox™ User's Guide

© COPYRIGHT 2015–2023 by The MathWorks, Inc.

Trademarks

MATLAB and Simulink are registered trademarks of The MathWorks, Inc. See www.mathworks.com/trademarks for a list of additional trademarks. Other product or brand names may be trademarks or registered trademarks of their respective holders.

Patents

MathWorks products are protected by one or more U.S. patents. Please see www.mathworks.com/patents for more information.

Revision History

Contents

Code Generation and Deployment

3

Streaming Pixel Interface

Streaming Pixel Interface

In this section...

What Is a Streaming Pixel Interface?

In hardware, processing an entire frame of video at one time has a high cost in memory and area. To save resources, serial processing is preferable in HDL designs. Vision HDL Toolbox blocks and System objects operate on a pixel, line, or neighborhood rather than a frame. The blocks and objects accept and generate video data as a serial stream of pixel data and control signals. The control signals indicate the relative location of each pixel within the image or video frame. The protocol mimics the timing of a video system, including inactive intervals between frames. Each block or object operates without full knowledge of the image format, and can tolerate imperfect timing of lines and frames.

All Vision HDL Toolbox blocks and System objects support single pixel streaming (with 1 pixel per cycle). Some blocks and System objects also support multipixel streaming (with 2, 4, or 8 pixels per cycle) for high-rate or high-resolution video. Multipixel streaming increases hardware resources to support higher video resolutions with the same hardware clock rate as a smaller resolution video. HDL code generation for multipixel streaming is not supported with System objects. Use the equivalent blocks to generate HDL code for multipixel algorithms.

How Does a Streaming Pixel Interface Work?

Video capture systems scan video signals from left to right and from top to bottom. As these systems scan, they generate inactive intervals between lines and frames of active video.

The *horizontal blanking* interval is made up of the inactive cycles between the end of one line and the beginning of the next line. This interval is often split into two parts: the *front porch* and the *back porch*. These terms come from the synchronize pulse between lines in analog video waveforms. The *front porch* is the number of samples between the end of the active line and the synchronize pulse. The *back porch* is the number of samples between the synchronize pulse and the start of the active line.

The *vertical blanking* interval is made up of the inactive cycles between the *ending active line* of one frame and the *starting active line* of the next frame.

The scanning pattern requires start and end signals for both horizontal and vertical directions. The Vision HDL Toolbox streaming pixel protocol includes the blanking intervals, and allows you to configure the size of the active and inactive frame.
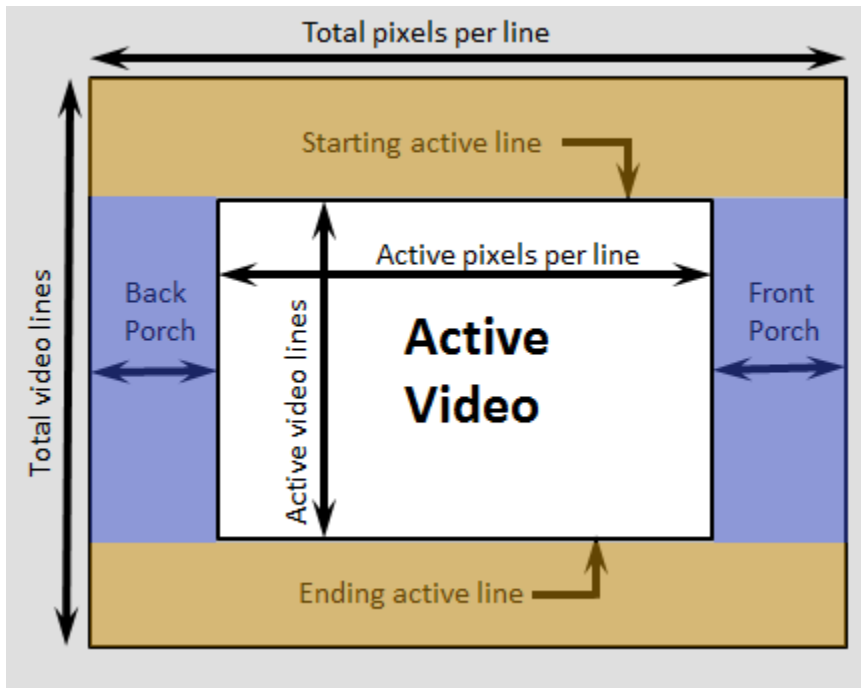
In the frame diagram, the blue shaded area to the left and right of the active frame indicates the horizontal blanking interval. The orange shaded area above and below the active frame indicates the vertical blanking interval. For more information on blanking intervals, see "Configure Blanking Intervals" on page 2-2.

Why Use a Streaming Pixel Interface?

Format Independence

The blocks and objects using this interface do not need a configuration option for the exact image size or the size of the inactive regions. In addition, if you change the image format for your design, you do not need to update each block or object. Instead, update the image parameters once at the serialization step. Some blocks and objects still require a line buffer size parameter to allocate memory resources.

By isolating the image format details, you can develop a design using a small image for faster simulation. Then once the design is correct, update to the actual image size.

Error Tolerance

Video can come from various sources such as cameras, tape storage, digital storage, or switching and insertion gear. These sources can introduce timing problems. Human vision cannot detect small variance in video signals, so the timing for a video system does not need to be perfect. Therefore, video processing blocks must tolerate variable timing of lines and frames.

By using a streaming pixel interface with control signals, each Vision HDL Toolbox block or object starts computation on a fresh segment of pixels at the start-of-line or start-of-frame signal. The computation occurs whether or not the block or object receives the end signal for the previous segment.

The protocol tolerates minor timing errors. If the number of valid and invalid cycles between start signals varies, the blocks or objects continue to operate correctly. Some Vision HDL Toolbox blocks and objects require minimum horizontal blanking regions to accommodate memory buffer operations. For more information, see "Configure Blanking Intervals" on page 2-2.

Pixel Stream Conversion Using Blocks and System Objects

In Simulink®, use the Frame To Pixels block to convert framed video data to a stream of pixels and control signals that conform to this protocol. The control signals are grouped in a nonvirtual bus data type called `pixelcontrol`. You can configure the block to return a pixel stream with 1, 2, 4, or 8 pixels per cycle.

In MATLAB®, use the `visionhdl.FrameToPixels` object to convert framed video data to a stream of pixels and control signals that conform to this protocol. The control signals are grouped in a structure data type. You can configure the object to create a pixel stream with 1, 2, 4, or 8 pixels per cycle.

If your input video is already in a serial format, you can design your own logic to generate `pixelcontrol` control signals from your existing serial control scheme. For example, see "Convert Camera Control Signals to pixelcontrol Format" on page 1-26 and "Integrate Vision HDL Blocks into Camera Link System" on page 1-31.

Supported Pixel Data Types

Vision HDL Toolbox blocks and objects include ports or arguments for streaming pixel data. Each block and object supports one or more pixel formats. The supported formats vary depending on the operation the block or object performs. This table details common video formats supported by Vision HDL Toolbox.

Type of Video	Pixel Format
Binary	Each pixel is represented by a single `boolean` or `logical` value. Used for true black-and-white video.
Grayscale	Each pixel is represented by *luma*, which is the gamma-corrected luminance value. This pixel is a single unsigned integer or fixed-point value.
Color	Each pixel is represented by 2 to 4 unsigned integer or fixed-point values representing the color components of the pixel. Vision HDL Toolbox blocks and objects use gamma-corrected color spaces, such as R'G'B' and Y'CbCr. To process multicomponent streams for blocks that do not support multicomponent input, replicate the block for each component. The `pixelcontrol` bus for all components is identical, so you can connect a single bus to multiple replicated blocks. To set up multipixel streaming for color video, you can configure the Frame To Pixels block to return a multicomponent and multipixel stream. See "MultiPixel-MultiComponent Video Streaming" on page 1-20.

Vision HDL Toolbox blocks have an input or output port, `pixel`, for the pixel data. Vision HDL Toolbox System objects expect or return an argument representing the pixel data. The following table describes the format of the pixel data.

Port or Argument	Description	Data Type
`pixel`	• Single pixel streaming — A scalar that represents a binary or grayscale pixel value or a row vector of two to four values representing a color pixel • Multipixel streaming — Column vector of two, four, or eight pixel values • Multipixel-multicomponent streaming — Matrix of two, four, or eight pixel values by two to four color components. You can simulate System objects with a multipixel streaming interface, but you cannot generate HDL code for System objects that use multipixel streams. To generate HDL code for multipixel algorithms, use the equivalent Simulink blocks.	Supported data types can include: • `boolean` or `logical` • `uint` or `int` • `fixdt()` The software supports `double` and `single` data types for simulation, but not for HDL code generation.

Note The blocks in this table support multipixel input, but not multicomponent pixels. The table shows what number of input pixels each block supports.

Block	Number of pixels
Image Filter	2, 4, or 8
Bilateral Filter	2, 4, or 8
Line Buffer	2, 4, or 8
Gamma Corrector	2, 4, or 8
Edge Detector	2, 4, or 8
Median Filter	2, 4, or 8
Histogram	2, 4, or 8
Lookup Table	2, 4, or 8
Binary morphology: Closing, Dilation, Erosion, and Opening	4 or 8

These blocks support multipixel-multicomponent pixel streams. The table shows what number of pixels and components each block supports.

Block	Number of pixels	Number of components
Pixel Stream FIFO	2, 4, or 8	1, 3, or 4
Color Space Converter	2, 4, or 8	3
Demosaic Interpolator	2, 4 or 8	3 (output only)
ROI Selector	2, 4, or 8	1, 3, or 4
Pixel Stream Aligner	2, 4, or 8	1, 3, or 4

Streaming Pixel Control Signals

Vision HDL Toolbox blocks and objects include ports or arguments for control signals relating to each pixel. These five control signals indicate the validity of a pixel and its location in the frame. For multipixel streaming, each vector of pixel values has one set of control signals.

In Simulink, the control signal port is a nonvirtual bus data type called `pixelcontrol`. For details of the bus data type, see "Pixel Control Bus" on page 1-24.

In MATLAB, the control signal argument is a structure. For details of the structure data type, see "Pixel Control Structure" on page 1-25.

Sample Time

Because the Frame To Pixels block creates a serial stream of the pixels of each input frame, the sample time of your video source must match the total number of pixels in the frame. The total number of pixels is *Total pixels per line × Total video lines*, so set the sample time to this value.

If your frame size is large, you may reach the fixed-step solver step size limit for sample times in Simulink, and receive an error like this.

```
The computed fixed step size (1.0) is 1000000.0 times smaller than all
the discrete sample times in the model.
```

You can avoid this error by choosing the variable-step solver.

Timing Diagram of Single Pixel Serial Interface

To illustrate the streaming pixel protocol, this example converts a frame to a sequence of control and data signals. Consider a 2-by-3 pixel image. To model the blanking intervals, configure the serialized image to include inactive pixels in these areas around the active image:

- 1-pixel-wide back porch
- 2-pixel-wide front porch
- 1 line before the first active line
- 1 line after the last active line

You can configure the dimensions of the active and inactive regions with the Frame To Pixels block or the `visionhdl.FrameToPixels` object.

In the figure, the active image area is in the dashed rectangle, and the inactive pixels surround it. The pixels are labeled with their grayscale values.
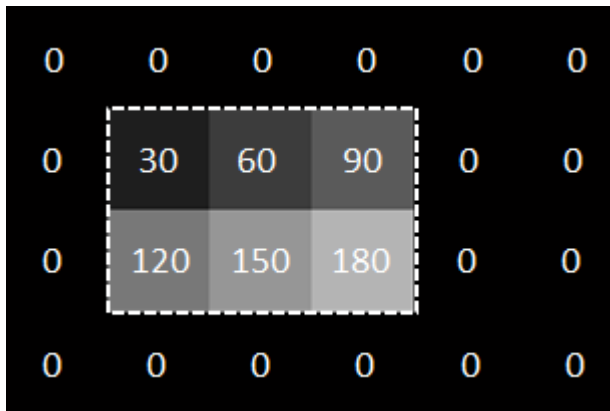
The block or object serializes the image from left to right, one line at a time. The timing diagram shows the control signals and pixel data that correspond to this image, which is the serial output of the Frame To Pixels block for this frame, configured for single-pixel streaming.
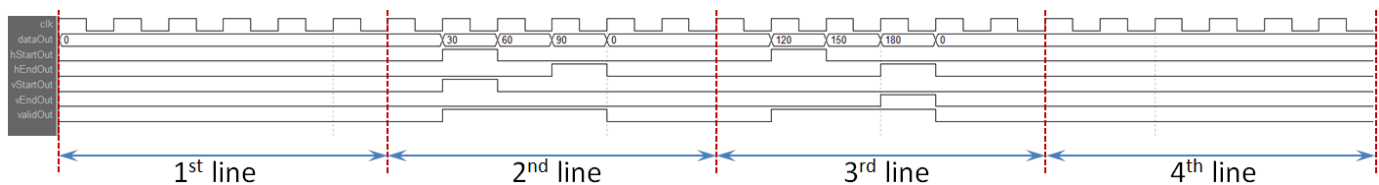
For an example using the Frame to Pixels block to serialize an image, see "Design Video Processing Algorithms for HDL in Simulink".

For an example using the `FrameToPixels` object to serialize an image, see "Design Hardware-Targeted Image Filters in MATLAB".

Timing Diagram of Multipixel Serial Interface

This example converts a frame to a multipixel stream with 4 pixels per cycle and corresponding control signals. Consider a 64-pixel-wide frame with these inactive areas around the active image.

- 4-pixel-wide back porch
- 4-pixel-wide front porch
- 4 lines before the first active line
- 4 lines after the last active line

The Frame to Pixels block configured for multipixel streaming returns pixel vectors formed from the pixels of each line in the frame from left to right. This diagram shows the top-left corner of the frame. The gray pixels show the active area of the frame, and the zero-value pixels represent blanking pixels. The label on each active pixel represents the location of the pixel in the frame. The highlighted boxes show the sets of pixels streamed on one cycle. The pixels in the inactive region are also streamed four at a time. The gray box shows the four blanking pixels streamed the cycle before the start of the active frame. The blue box shows the four pixel values streamed on the first valid cycle of the frame, and the orange box shows the four pixel values streamed on the second valid cycle of the frame. The green box shows the first four pixels of the next active line.

This waveform shows the multipixel streaming data and control signals for the first line of the same frame, streamed with 4 pixels per cycle. The `pixelcontrol` signals that apply to each set of four pixel values are shown below the data signals. Because the vector has only one `valid` signal, the pixels in the vector are either all valid or all invalid. The `hStart` and `vStart` signals apply to the pixel with the lowest index in the vector. The `hEnd` and `vEnd` signals apply to the pixel with the highest index in the vector.

Prior to the time period shown, the initial vertical blanking pixels are streamed four at a time, with all control signals set to `false`. This waveform shows the pixel stream of the first line of the image. The gray, blue, and orange boxes correspond to the highlighted areas of the frame diagram. After the first line is complete, the stream has two cycles of horizontal blanking that contains 8 invalid pixels (front and back porch). Then, the waveform shows the next line in the stream, starting with the green box.
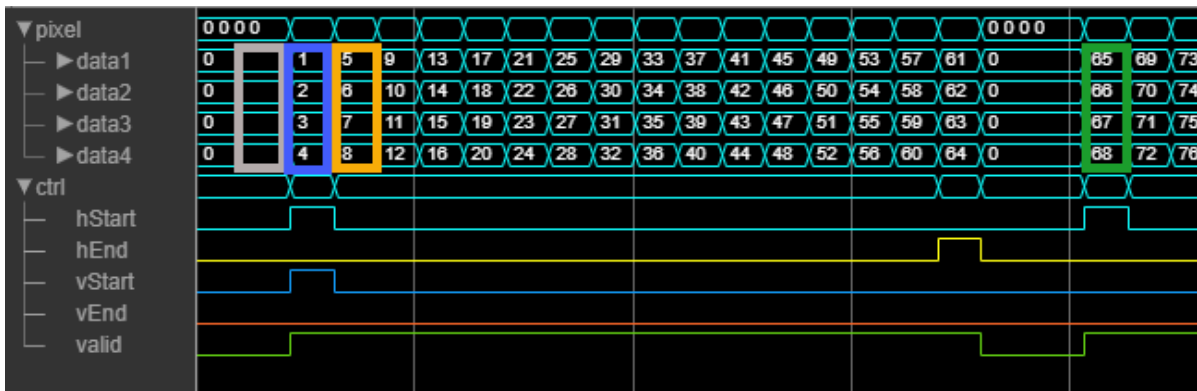
For an example model that uses multipixel streaming, see "Filter Multipixel Video Streams" on page 1-10.

See Also

Frame To Pixels | Pixels To Frame | `visionhdl.FrameToPixels` | `visionhdl.PixelsToFrame`

Related Examples

- "Design Video Processing Algorithms for HDL in Simulink"
- "Design Hardware-Targeted Image Filters in MATLAB"
- "Filter Multipixel Video Streams" on page 1-10

- "MultiPixel-MultiComponent Video Streaming" on page 1-20

Filter Multipixel Video Streams

This example shows how to design filters that operate on a multipixel input video stream. Use multipixel streaming to process high-resolution or high-frame-rate video with the same synthesized clock frequency as a single-pixel streaming interface. Multipixel streaming also improves simulation speed and throughput because fewer iterations are required to process each frame, while maintaining the hardware benefits of a streaming interface.

The example model has three subsystems which each perform the same algorithm:

- **SinglePixelGaussianEdge**: Uses the Image Filter and Edge Detector blocks to operate on a single-pixel stream. This subsystem shows how the rates and interfaces for single-pixel streaming compare with multipixel designs.
- **MultiPixelGaussianEdge**: Uses the Image Filter and Edge Detector blocks to operate on a multipixel stream. This subsystem shows how to use the multipixel interface with library blocks.
- **MultiPixelCustomGaussianEdge**: Uses the Line Buffer block to build a Gaussian filter and Sobel edge detection for a multipixel stream. This subsystem shows how to use the Line Buffer output for multipixel design.

Processing multipixel video streams allows for higher frame rates to be achieved without a corresponding increase to the clock frequency. Each of the subsystems can achieve 200MHz clock frequency on a Xilinx ZC706 board. The 480p video stream has **Total pixels per line** x **Total video lines** = 800*525 cycles per frame. With a single pixel stream you can process 200M/(800*525) = 475 frames per second. In the multipixel subsystem, 4 pixels are processed on each cycle, which reduces the number of cycles per line to 200. This means that with a multipixel stream operating on 4 pixels at a time, at 200MHz, on a 480p stream, 1900 frames can be processed per second. If the resolution is increased from 480p to 1080p, 80 frames per second can be achieved in the single pixel case versus 323 frames per second for 4 pixels at a time or 646 frames per second for 8 pixels at a time.
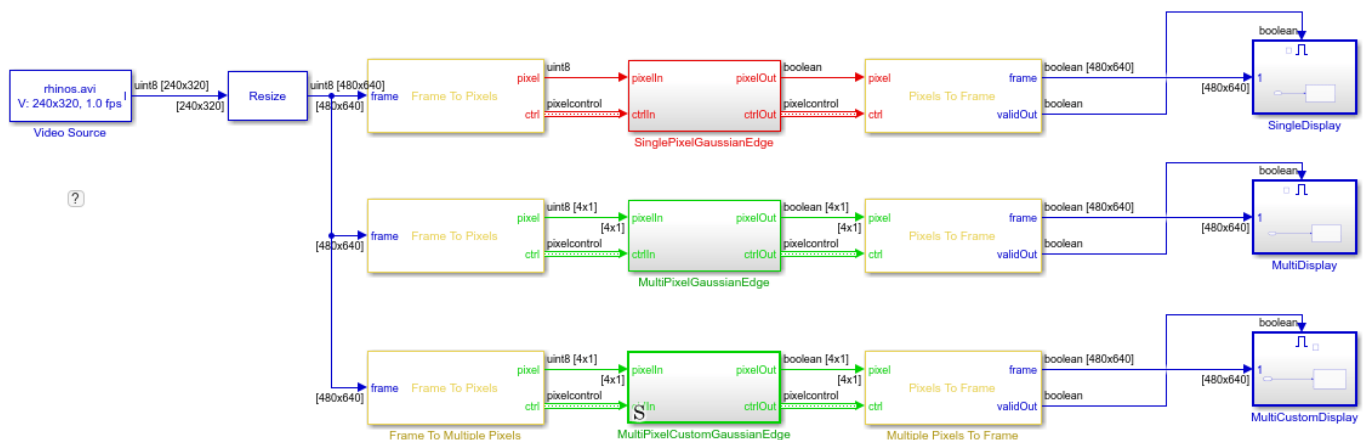
Copyright 2019 The MathWorks, Inc.

Multipixel Streaming Using Library Blocks

Generate a multipixel stream from the Frame to Pixels block by setting **Number of pixels** to 4 or 8. The default value of 1 returns a scalar pixel stream with a sample rate of **Total pixels per line** * **Total video lines** faster than the frame rate. This rate shows red in the example model. The two

multipixel subsystems use a multipixel stream with **Number of pixels** set to 4. This configuration returns 4 pixels on each clock cycle and has a sample rate of (**Total pixels per line**/4) * **Total video lines**. The lower output rate, which is green in the model, shows that you can increase either the input frame rate or resolution by a factor of 4 and therefore process 4 times as many pixels in the same frame period using the same clock frequency as the single pixel case.

The **SinglePixelGaussianEdge** and **MultiPixelGaussianEdge** subsystems compute the same result using the Image Filter and Edge Detector blocks.

In **MultiPixelGaussianEdge**, the blocks accept and return four pixels on each clock cycle. You do not have to configure the blocks for multipixel streaming, they detect the input size on the port. The `pixelcontrol` bus indicates the validity and location in the frame of each set of four pixels. The blocks buffer the [4x1] stream to form four [*KernelHeight* x *KernelWidth*] kernels, and compute four convolutions in parallel to give a [4x1] output.
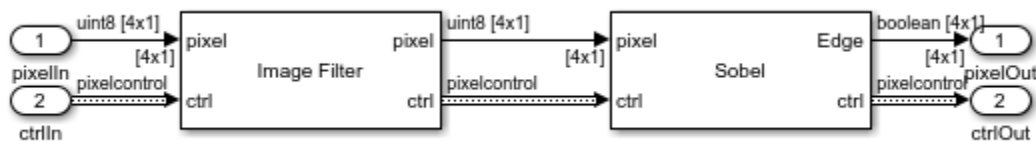
Custom Multipixel Algorithms

The **MultiPixelCustomGaussianEdge** subsystem uses the Line Buffer block to implement a custom filtering algorithm. This subsystem is similar to how the library blocks internally implement multipixel kernel operations. The Image Filter and Edge Detector blocks use more detailed optimizations than are shown here. This implementation shows a starting point for building custom multipixel algorithms using the output of the Line Buffer block.

The custom filter and custom edge detector use the Line Buffer block to return successive [*KernelHeight* x *NumberofPixels*] regions. Each region is passed to the KernelIndexer subsystem which uses buffering and indexing logic to form Number of Pixels * [*KernelHeight* x *KernelWidth*] filter kernels. Then each kernel is passed to a separate FilterKernel subsystem to perform convolutions in parallel.
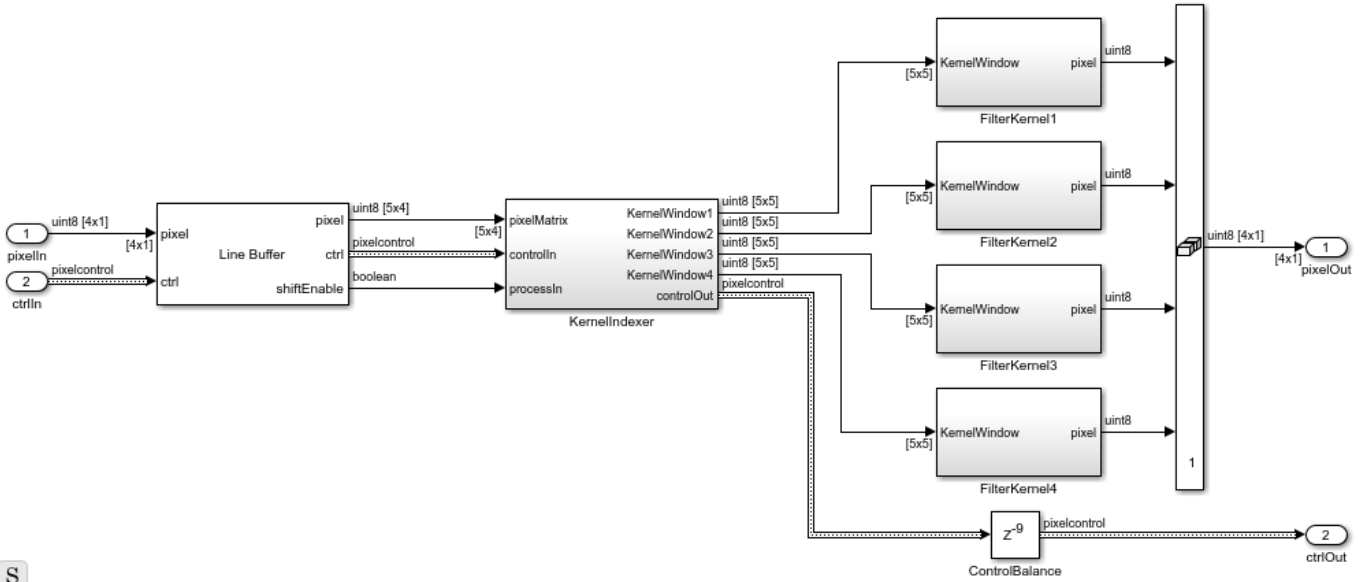
Form Kernels from Line Buffer Output

The KernelIndexer subsystem forms 4 [5x5] filter kernels from the 2-D output of the Line Buffer block.
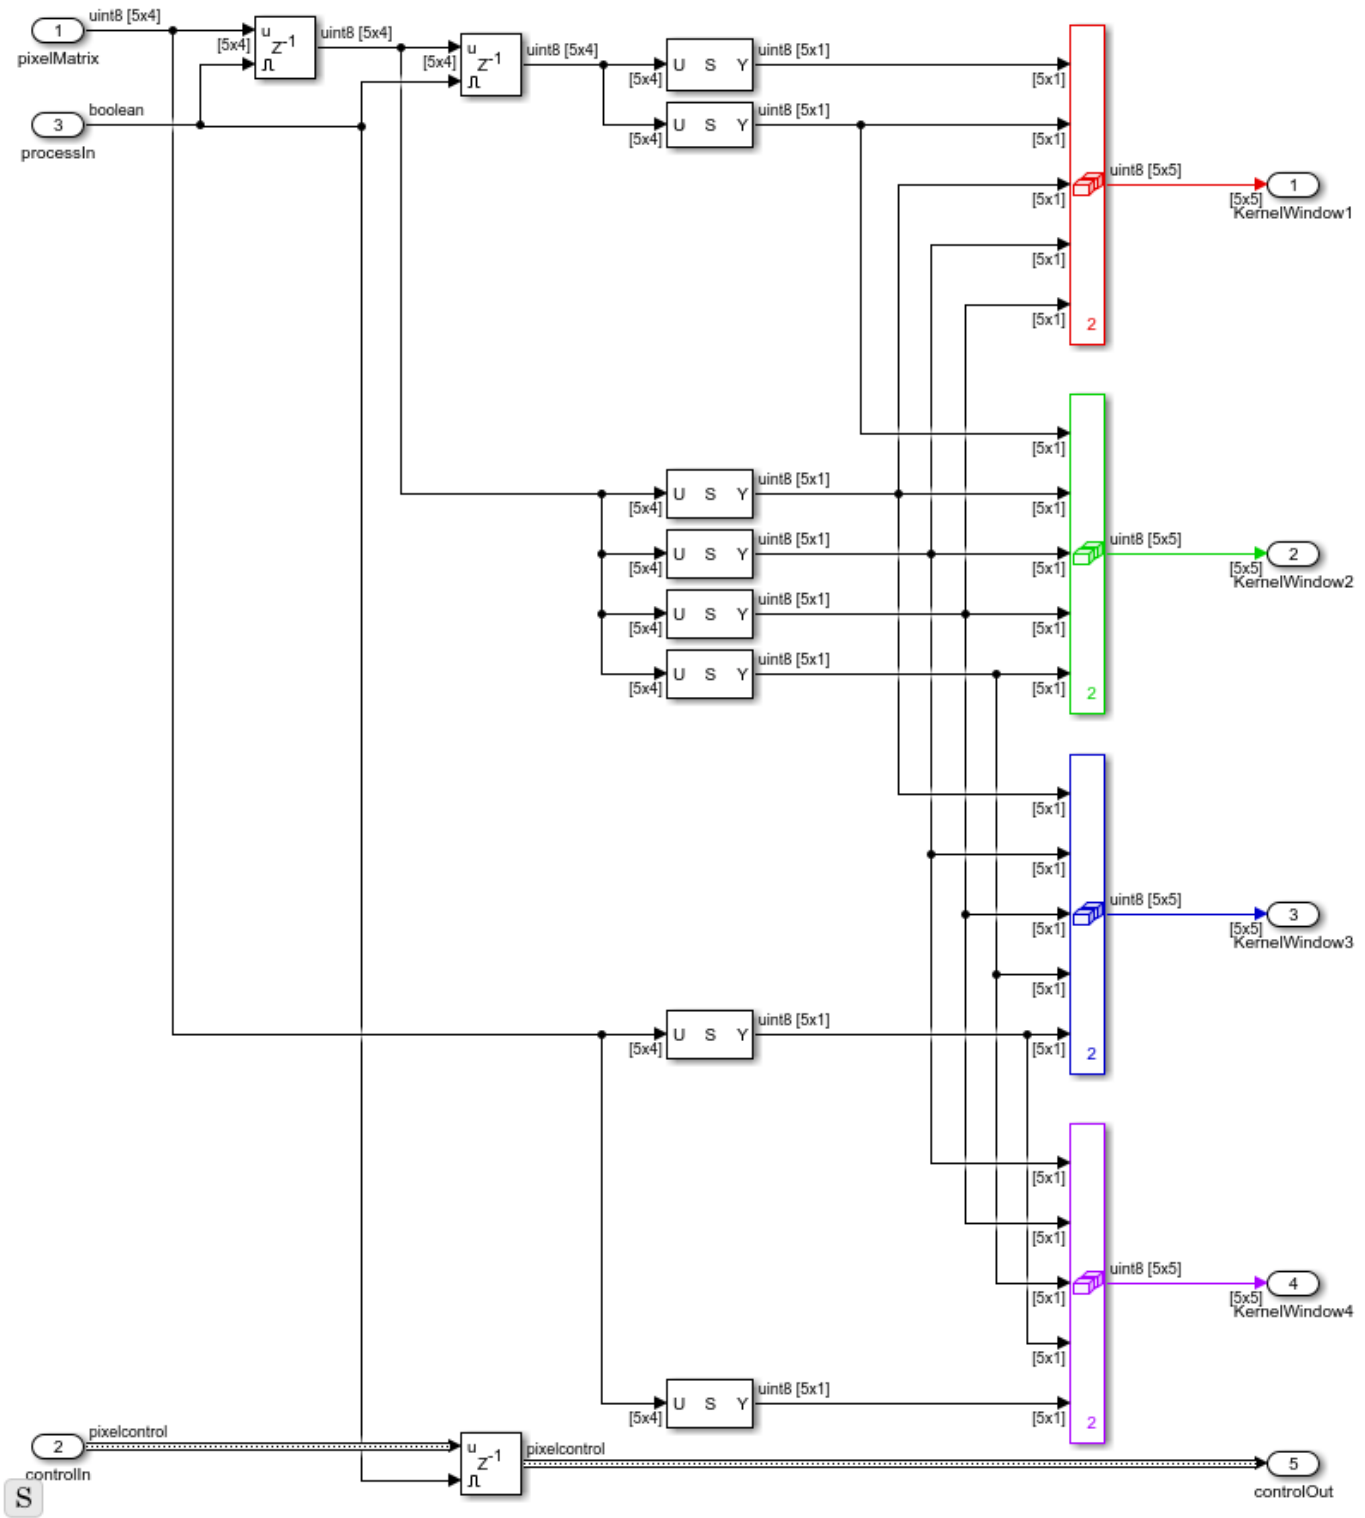
The diagram shows how the filter kernel is extracted from the [5x4] output stream, for the kernel that is centered on the first pixel in the [4x1] output. This first kernel includes pixels from 2 adjacent [5x4] Line Buffer outputs.

P	P	P	P	P	P	P	P	P	P	P
P	P	P	P	P	P	P	P	P	P	P
P	P	1	2	3	4	5	6	7	8	9
P	P	33	34	35	36	37	38	39	40	41
P	P	65	66	67	68	69	70	71	72	73
P	P	97	98	99	100	101	102	103	104	105
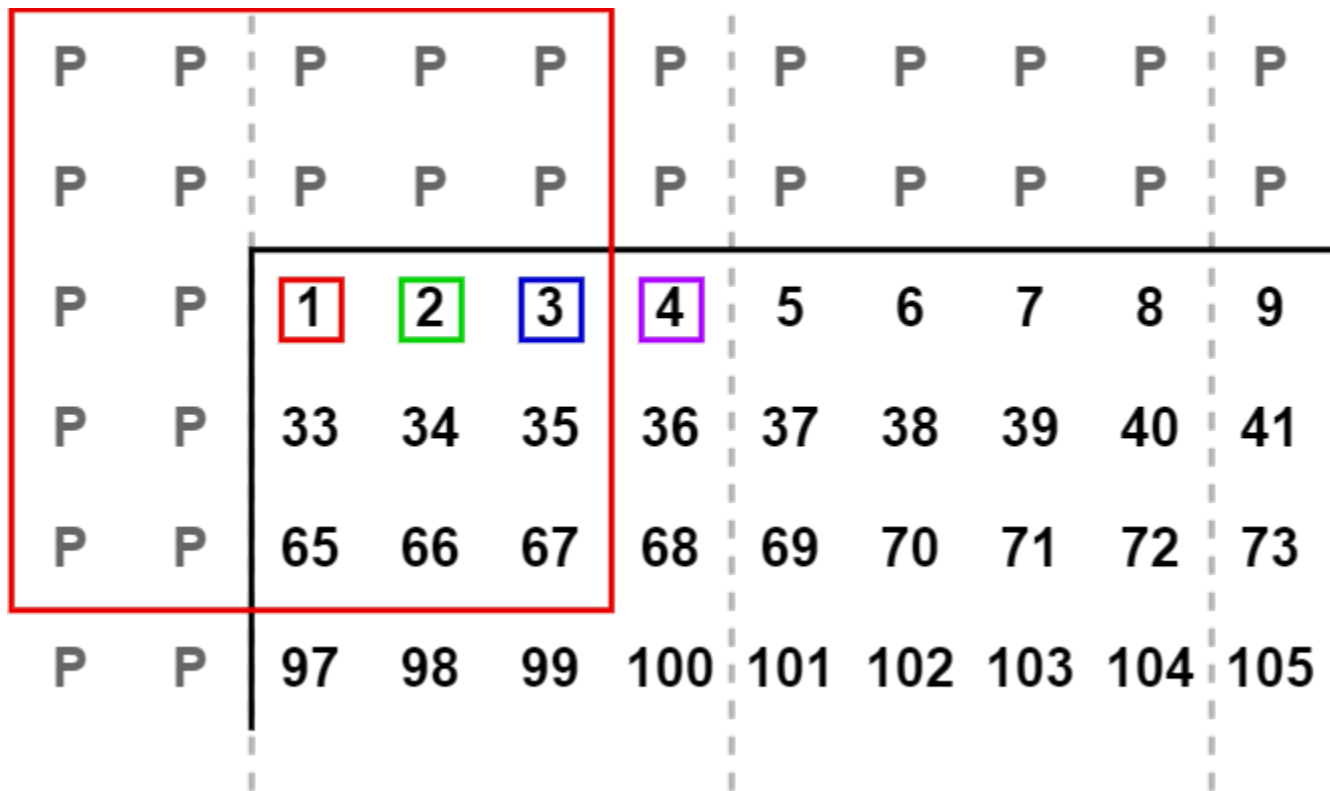
The kernel centered on the last pixel in the [4x1] output also includes the third adjacent [5x4] output. So, to form four [5x5] kernels, the subsystem must access columns from three [5x4] matrices.
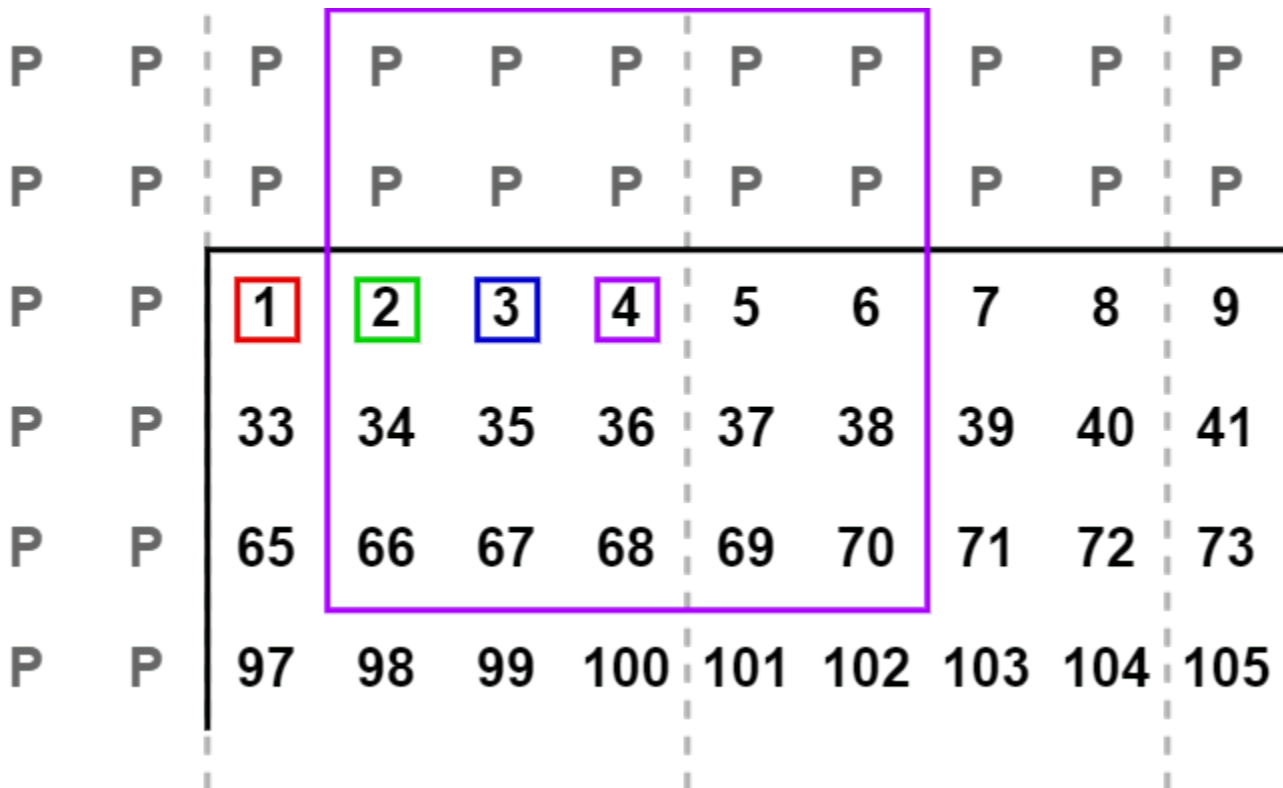
The KernelIndexer subsystem uses the current [5x4] input, and stores two more [5x4] matrices using registers enabled by `shiftEnable`. This design is similar to the tapped delay line used with a Line Buffer using single pixel streaming. The subsystem then accesses pixel data across the columns to form the four [5x5] kernels. The Image Filter block uses this same logic internally when the block has multipixel input. The block automatically designs this logic at compile time for any supported kernel size.

Implement Filters

Since the input multipixel stream is a [4x1] vector, the filters must perform four convolutions on each cycle to keep pace with the incoming data. There are four parallel FilterKernel subsystems that each perform the same operation. The [5x5] matrix multiply is implemented as a [5x1] vector multiply by using a For Each subsystem containing a pipelined multiplier. The output is passed to an adder tree. The adder tree is also pipelined, and the pipeline latency is applied to the `pixelcontrol` signal to match. The results of the four FilterKernel subsystems are then concatenated into a [4x1] output vector.
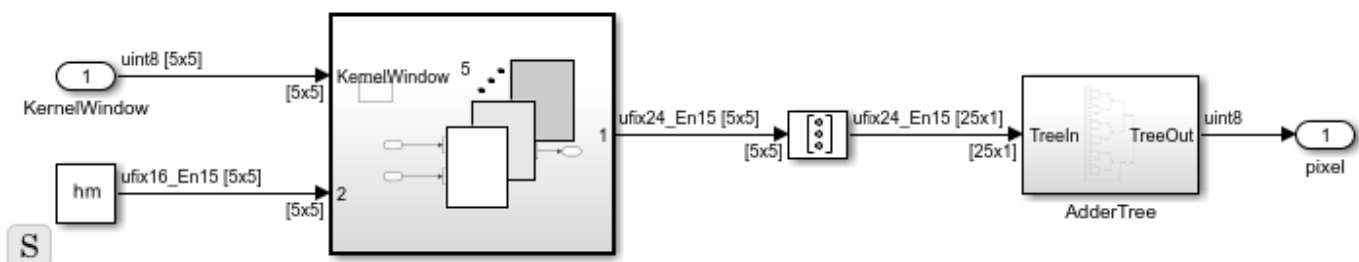
Implement Edge Detectors

To match the algorithm of the Edge Detector block, this custom edge detector uses a [3x3] kernel size. Compare this KernelIndexer subsystem for the [3x3] edge detection with the [5x5] kernel described above. The algorithm still must access three successive matrices from the output of the Line Buffer block (including padding on either side of the kernel). However, the algorithm saves fewer columns to form a smaller filter kernel.
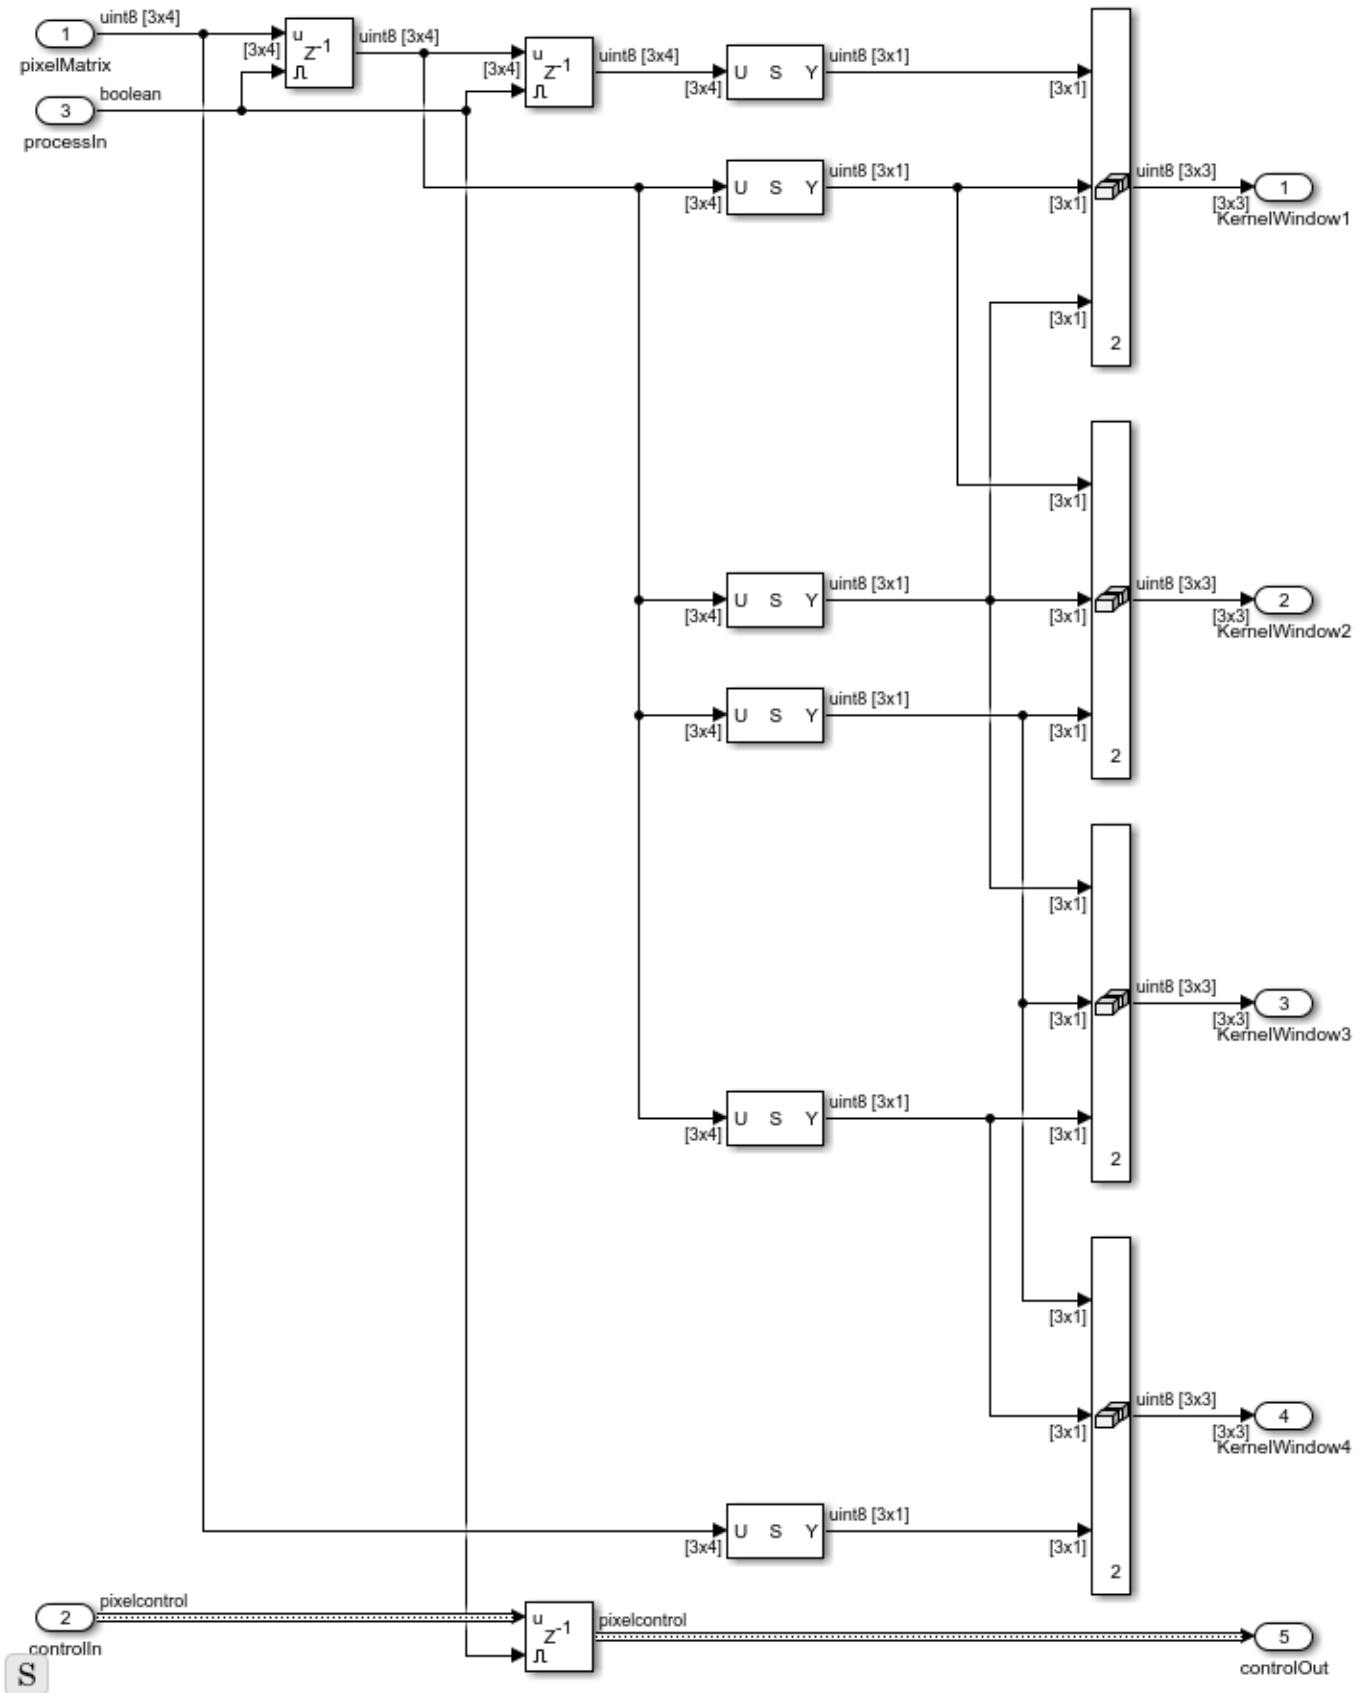
Extending to Larger Kernel Sizes

For larger kernel sizes the number of [*KernelHeight* x *NumPixels*] regions to store in the KernelIndexer is `(2 * ceil(floor(`*KernelWidth* `/ 2) /` *NumPixels*`) + 1)`. In such a case, the number of inputs to the concatenators increases to *KernelWidth* and you must route these additional inputs from the tapped delay line of Line Buffer matrices. For a [4x1] multipixel stream with a [11x11] kernel size you would need to store five [11x4] matrices from the Line Buffer to form four [11x11] kernels each cycle.

P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P
P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P
P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P
P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P
P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P	P
P	P	P	P	P	1	2	3	4	5	6	7	8	9	10	11	12	13
P	P	P	P	P	33	34	35	36	37	38	39	40	41	42	43	44	45
P	P	P	P	P	65	66	67	68	69	70	71	72	73	74	75	76	77
P	P	P	P	P	97	98	99	100	101	102	103	104	105	106	107	108	109
P	P	P	P	P	129	130	131	132	133	134	135	136	137	138	139	140	141
P	P	P	P	P	161	162	163	164	165	166	167	168	169	170	171	172	173
P	P	P	P	P	161	162	163	164	165	166	167	168	169	170	171	172	173

Improving Simulation Time

In the default example configuration, the single pixel, multipixel, and custom multipixel subsystems all run in parallel. The simulation speed is limited by the time processing the single-pixel path because it requires more iterations to process the same size of frame. To observe the simulation speed improvement for multipixel streaming, comment out the single-pixel data path.

HDL Implementation Results

HDL was generated from both the **MultiPixelGaussianEdge** subsystem and the **MultiPixelCustomGaussianEdge** subsystem and put through Place and Route on a Xilinx™ ZC706 board. The **MultiPixelCustomGaussianEdge** subsystem, which does not attempt to optimize coefficients, had the following results -

```
T =

  4x2 table

    Resource      Usage
    _____      _____

    DSP48         108
    Flip Flop     9842
    LUT           4960
    BRAM          12
```

The **MultiPixelGaussianEdge** subsystem, which uses the optimized Image Filter and Edge Detector blocks uses less resources, as shown in the table below. This comparison shows the resource savings achieved because the blocks analyze the filter structure and pre-add repeated coefficients.

```
T =

  4x2 table

    Resource      Usage
    _____      _____

    DSP48         16
    Flip Flop     3959
    LUT           1789
    BRAM          10
```

See Also
Frame To Pixels | Image Filter | Edge Detector | Pixels To Frame

More About
• "Streaming Pixel Interface" on page 1-2

MultiPixel-MultiComponent Video Streaming

This example shows how to work with a multipixel-multicomponent pixel stream. Multipixel-multicomponent streaming enables real-time processing of high-resolution or high-frame-rate color video streams.

To demonstrate working with such a video stream, this example implements the well-known *bloom effect* image post-processing technique. The bloom effect introduces or enhances the glow of light sources in an image.

Top Level I/0

Each pixel of a high-resolution or high-frame-rate pixel stream is modeled as a NumPixel-by-NumComponent matrix. Matrix data types are supported for HDL code generation within a design, but not for the ports of the top-level subsystem. In this case, the input pixel stream is split into three 4-by-1 vectors at the input of the DUT, and then recombined at the output into a 4-by-3 matrix for the Pixels To Frame block.
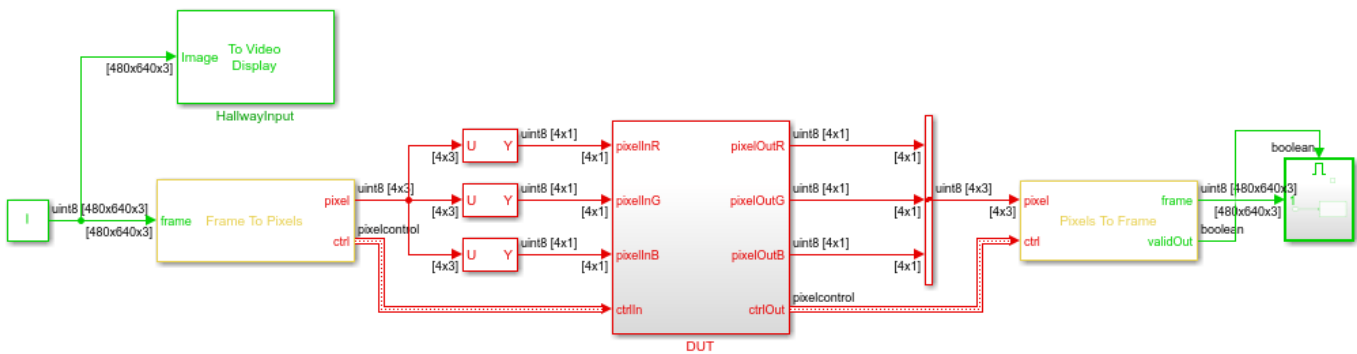
Copyright 2019-2022 The MathWorks, Inc.

Bloom Effect

The example model follows these three steps to add a bloom effect to the input image.

1 The BrightSpotFilter subsystems find bright spots in the intensity image by checking pixel values against a threshold. These pixel values have been converted from RGB to intensity by the Color Space Converter.

2 The 15x15 Blur subsystem spreads out the bright spots by applying a Gaussian filter.

3 The BlendBloom subsystem adds the Gaussian-enhanced bright spots back to the original image.
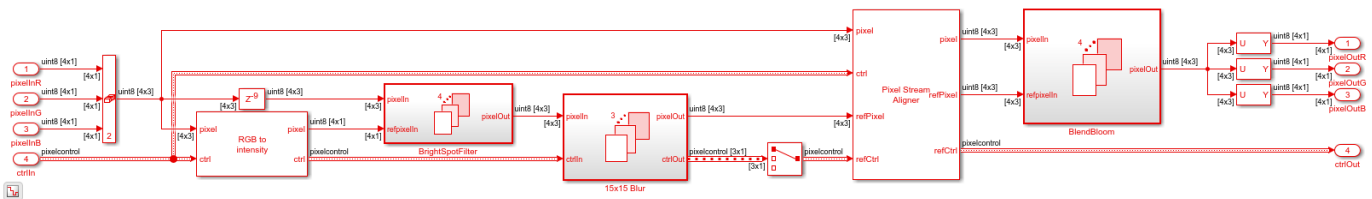
Matrix Operations

Vision HDL Toolbox™ neighborhood-processing blocks can operate on vector inputs, but do not support matrix inputs. The line buffer used inside the blocks returns a NumPixels-by-KernelHeight matrix. Using multicomponent inputs would result in a NumPixels-by-KernelHeight-by-NumComponents output matrix, however, 3-D matrices are not supported for HDL code generation. To work around this the model uses For Each subsystems which support HDL code generation with scalar, vector, and matrix inputs.

The model sets the **Partition Width** and **Partition Dimension** of the BrightSpotFilter to 1. The subsystem applies a threshold in parallel to each 1x3 RGB multi-component pixel of its input multipixel-multicomponent matrix.

The 15x15 Blur subsystem has the **Partition Width** set to 1 and **Partition Dimension** set to 2. The subsystem applies a Gaussian filter in parallel to each 4x1 multipixel RGB component.

The BlendBloom subsystem has the **Partition Width** and **Partition Dimension** set to 1. The subsystem adds in parallel each 1x3 multicomponent pixel to its respective filtered 1x3 multicomponent pixel.

Simulation Results

Simulating the model displays these input and output images. The bloom effect makes the lighted areas of the scene look brighter and shows a halo effect.
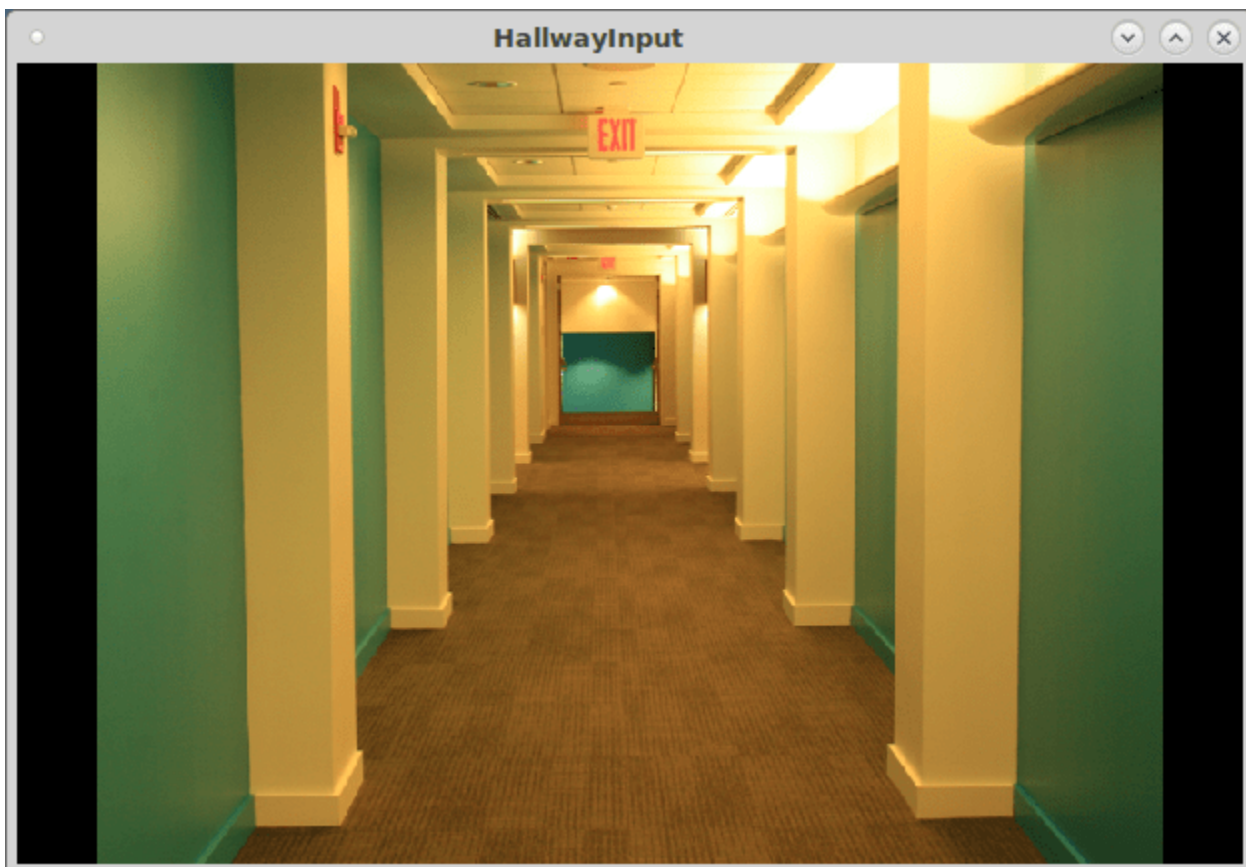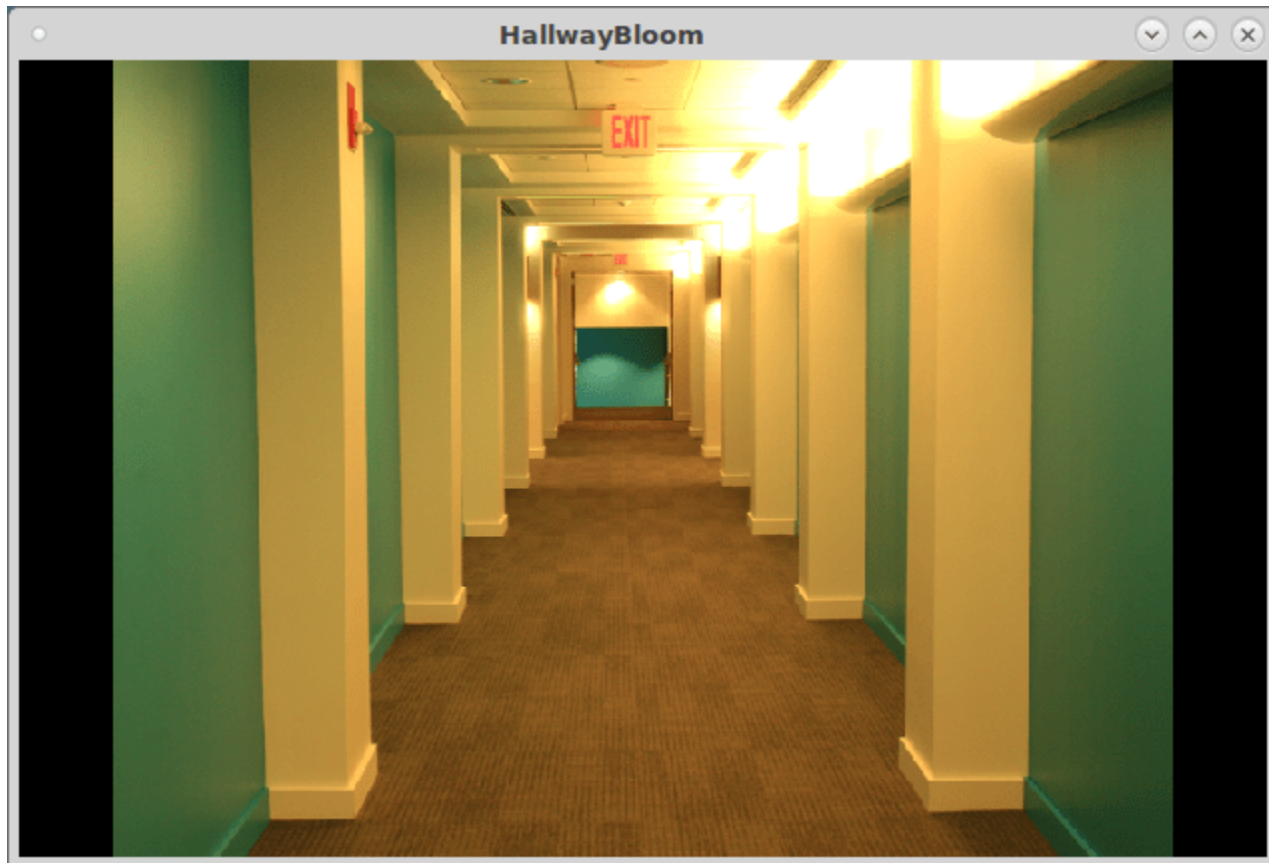
Implementation Results

This table shows the synthesis results of HDL code generated from the DUT subsystem and synthesized for a Xilinx™ Zynq™ ZC706 board. Because none of the resources exceed 25% of their respective category, the design has a relatively small footprint.

```
T =

  4x2 table

    Resource     Usage
    _____    _____

    DSP48        84
    Flip Flop    61739
    LUT          36966
    BRAM         132
```

See Also

Frame To Pixels | Pixels To Frame

More About

- "Streaming Pixel Interface" on page 1-2

Pixel Control Bus

Vision HDL Toolbox blocks use a nonvirtual bus data type, `pixelcontrol`, for control signals associated with serial pixel data. The bus contains 5 `boolean` signals indicating the validity of a pixel and its location within a frame. You can easily connect the data and control output of one block to the input of another, because Vision HDL Toolbox blocks use this bus for input and output. To convert an image into a pixel stream and a `pixelcontrol` bus, use the Frame to Pixels block.

Signal	Description	Data Type
hStart	`true` for the first pixel in a horizontal line of a frame	boolean
hEnd	`true` for the last pixel in a horizontal line of a frame	boolean
vStart	`true` for the first pixel in the first (top) line of a frame	boolean
vEnd	`true` for the last pixel in the last (bottom) line of a frame	boolean
valid	`true` for any valid pixel	boolean

For multipixel streaming, each vector of pixel values has one set of control signals. Because the vector has only one `valid` signal, the pixels in the vector must be either all valid or all invalid. The `hStart` and `vStart` signals apply to the pixel with the lowest index in the vector. The `hEnd` and `vEnd` signals apply to the pixel with the highest index in the vector.

Troubleshooting: When you generate HDL code from a Simulink model that uses this bus, you may need to declare an instance of `pixelcontrol` bus in the base workspace. If you encounter the error `Cannot resolve variable 'pixelcontrol'` when you generate HDL code in Simulink, use the `pixelcontrolbus` function to create an instance of the bus type. Then try generating HDL code again.

To avoid this issue, the Vision HDL Toolbox model template includes this line in the `InitFcn` callback.

```
evalin('base','pixelcontrolbus')
```

See Also
Frame To Pixels | Pixels To Frame | `pixelcontrolbus`

More About
- "Streaming Pixel Interface" on page 1-2

Pixel Control Structure

Vision HDL Toolbox System objects use a structure data type for control signals associated with serial pixel data. The structure contains five `logical` signals indicating the validity of a pixel and its location within a frame. You can easily pass the data and control output arguments of one Vision HDL Toolbox System object™ as the input arguments to another Vision HDL Toolbox System object, because the objects use this structure for input and output control signal arguments. To convert an image into a pixel stream and control signals, use the `visionhdl.FrameToPixels` System object.

Signal	Description	Data Type
hStart	`true` for the first pixel in a horizontal line of a frame	`logical`
hEnd	`true` for the last pixel in a horizontal line of a frame	`logical`
vStart	`true` for the first pixel in the first (top) line of a frame	`logical`
vEnd	`true` for the last pixel in the last (bottom) line of a frame	`logical`
valid	`true` for any valid pixel	`logical`

See Also
`pixelcontrolstruct` | `pixelcontrolsignals` | `visionhdl.FrameToPixels` | `visionhdl.PixelsToFrame`

More About
*   "Streaming Pixel Interface" on page 1-2

Convert Camera Control Signals to pixelcontrol Format

This example shows how to convert Camera Link® signals to the `pixelcontrol` structure, invert the pixels with a Vision HDL Toolbox™ object, and convert the control signals back to the Camera Link format.

Vision HDL Toolbox blocks and objects use a custom streaming video format. If your system operates on streaming video data from a camera, you must convert the camera control signals into this custom format. Alternatively, if you integrate Vision HDL Toolbox algorithms into existing design and verification code that operates in the camera format, you must also convert the output signals from the Vision HDL Toolbox design back to the camera format.

You can generate HDL code from the three functions in this example. To create local copies of all the files in this example, so you can view and edit them, click the Open Script button.

Create Input Data in Camera Link Format

The Camera Link format consists of three control signals: F indicates the valid frame, L indicates each valid line, and D indicates each valid pixel. For this example, create input vectors in the Camera Link format to represent a basic padded video frame. The vectors describe this 2-by-3, 8-bit grayscale frame. In the figure, the active image area is in the dashed rectangle, and the inactive pixels surround it. The pixels are labeled with their grayscale values.
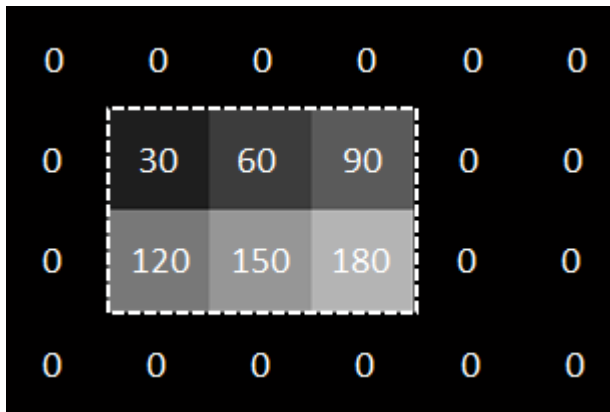

```
F = logical([0,0,0,0,0,0,0,1,1,1,1,1,1,1,1,1,0,0,0,0,0,0,0,0]);
L = logical([0,0,0,0,0,0,0,1,1,1,0,0,0,1,1,1,0,0,0,0,0,0,0,0]);
D = logical([0,0,0,0,0,0,0,1,1,1,0,0,0,1,1,1,0,0,0,0,0,0,0,0]);
pixel = uint8([0,0,0,0,0,0,0,30,60,90,0,0,0,120,150,180,0,0,0,0,0,0,0,0]);
```

Design Vision HDL Toolbox Algorithm

Create a function to invert the image using Vision HDL Toolbox algorithms. The function contains a System object that supports HDL code generation. This function expects and returns a pixel and associated control signals in Vision HDL Toolbox format.

```
function [pixOut,ctrlOut] = Invert(pixIn,ctrlIn)

persistent invertI;
if isempty(invertI)
```

```
    tabledata = linspace(255,0,256);
    invertI = visionhdl.LookupTable(uint8(tabledata));
end

[pixOut,ctrlOut] = invertI(pixIn,ctrlIn);
```

Convert Camera Link Control Signals to `pixelcontrol` Format

Write a custom System object to convert Camera Link signals to the Vision HDL Toolbox control signal format. The object converts the control signals, and then calls the `pixelcontrolstruct` function to create the structure expected by the Vision HDL Toolbox System objects. This code snippet shows the logic to convert the signals.

```
ctrl = pixelcontrolstruct(obj.hStartOutReg,obj.hEndOutReg,...
                    obj.vStartOutReg,obj.vEndOutReg,obj.validOutReg);

vStart = obj.FReg && ~obj.FPrevReg;
vEnd = ~F && obj.FReg;
hStart = obj.LReg && ~obj.LPrevReg;
hEnd = ~L && obj.LReg;

obj.vStartOutReg = vStart;
obj.vEndOutReg = vEnd;
obj.hStartOutReg = hStart;
obj.hEndOutReg = hEnd;
obj.validOutReg = obj.DReg;
```

The object stores the input and output control signal values in registers. `vStart` goes high for one cycle at the start of `F`. `vEnd` goes high for one cycle at the end of `F`. `hStart` and `hEnd` are derived similarly from `L`. The object returns the current value of `ctrl` each time you call it.

This processing adds two cycles of delay to the control signals. The object passes through the pixel value after matching delay cycles. For the complete code for the System object, see `CAMERALINKtoVHT_Adapter.m`.

Convert `pixelcontrol` to Camera Link

Write a custom System object to convert Vision HDL Toolbox signals back to the Camera Link format. The object calls the `pixelcontrolsignals` function to flatten the control structure into its component signals. Then it computes the equivalent Camera Link signals. This code snippet shows the logic to convert the signals.

```
[hStart,hEnd,vStart,vEnd,valid] = pixelcontrolsignals(ctrl);

Fnew = (~obj.FOutReg && vStart) || (obj.FPrevReg && ~obj.vEndReg);
Lnew = (~obj.LOutReg && hStart) || (obj.LPrevReg && ~obj.hEndReg);

obj.FOutReg = Fnew;
obj.LOutReg = Lnew;
obj.DOutReg = valid;
```

The object stores the input and output control signal values in registers. `F` is high from `vStart` to `vEnd`. `L` is high from `hStart` to `hEnd`. The object returns the current values of `FOutReg`, `LOutReg`, and `DOutReg` each time you call it.

This processing adds one cycle of delay to the control signals. The object passes through the pixel value after a matching delay cycle. For the complete code for the System object, see `VHTtoCAMERALINKAdapter.m`.

Create Conversion Functions That Support HDL Code Generation

Wrap the converter System objects in functions, similar to `Invert`, so you can generate HDL code for these algorithms.

```
function [ctrl,pixelOut] = CameraLinkToVisionHDL(F,L,D,pixel)
% CameraLink2VisionHDL : converts one cycle of CameraLink control signals
% to Vision HDL format, using a custom System object.
% Introduces two cycles of delay to both ctrl signals and pixel data.

persistent CL2VHT;
if isempty(CL2VHT)
    CL2VHT = CAMERALINKtoVHT_Adapter();
end

[ctrl,pixelOut] = CL2VHT(F,L,D,pixel);
```

See `CameraLinkToVisionHDL.m`, and `VisionHDLToCameraLink.m`.

Write a Test Bench

To invert a Camera Link pixel stream using these components, write a test bench script that:

1 Preallocates output vectors to reduce simulation time
2 Converts the Camera Link control signals for each pixel to the Vision HDL Toolbox format
3 Calls the `Invert` function to flip each pixel value
4 Converts the control signals for that pixel back to the Camera Link format

```
[~,numPixelsPerFrame] = size(pixel);
pixOut = zeros(numPixelsPerFrame,1,'uint8');
pixel_d = zeros(numPixelsPerFrame,1,'uint8');
pixOut_d = zeros(numPixelsPerFrame,1,'uint8');
DOut = false(numPixelsPerFrame,1);
FOut = false(numPixelsPerFrame,1);
LOut = false(numPixelsPerFrame,1);
ctrl = repmat(pixelcontrolstruct,numPixelsPerFrame,1);
ctrlOut = repmat(pixelcontrolstruct,numPixelsPerFrame,1);

for p = 1:numPixelsPerFrame
  [pixel_d(p),ctrl(p)] = CameraLinkToVisionHDL(pixel(p),F(p),L(p),D(p));
  [pixOut(p),ctrlOut(p)] = Invert(pixel_d(p),ctrl(p));
  [pixOut_d(p),FOut(p),LOut(p),DOut(p)] = VisionHDLToCameraLink(pixOut(p),ctrlOut(p));
end
```

View Results

The resulting vectors represent this inverted 2-by-3, 8-bit grayscale frame. In the figure, the active image area is in the dashed rectangle, and the inactive pixels surround it. The pixels are labeled with their grayscale values.
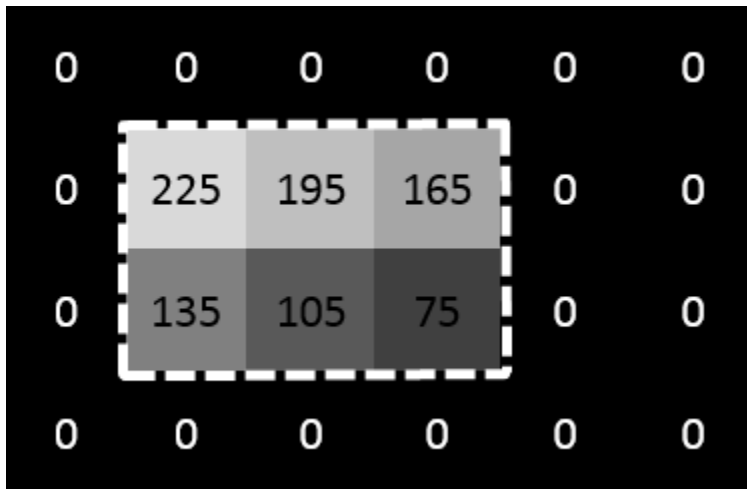
If you have a DSP System Toolbox™ license, you can view the vectors as signals over time using the Logic Analyzer. This waveform shows the `pixelcontrol` and Camera Link control signals, the starting pixel values, and the delayed pixel values after each operation. Run NewCustomCtrlSignals_LogicAnalyzer.m to generate this logic analyzer.
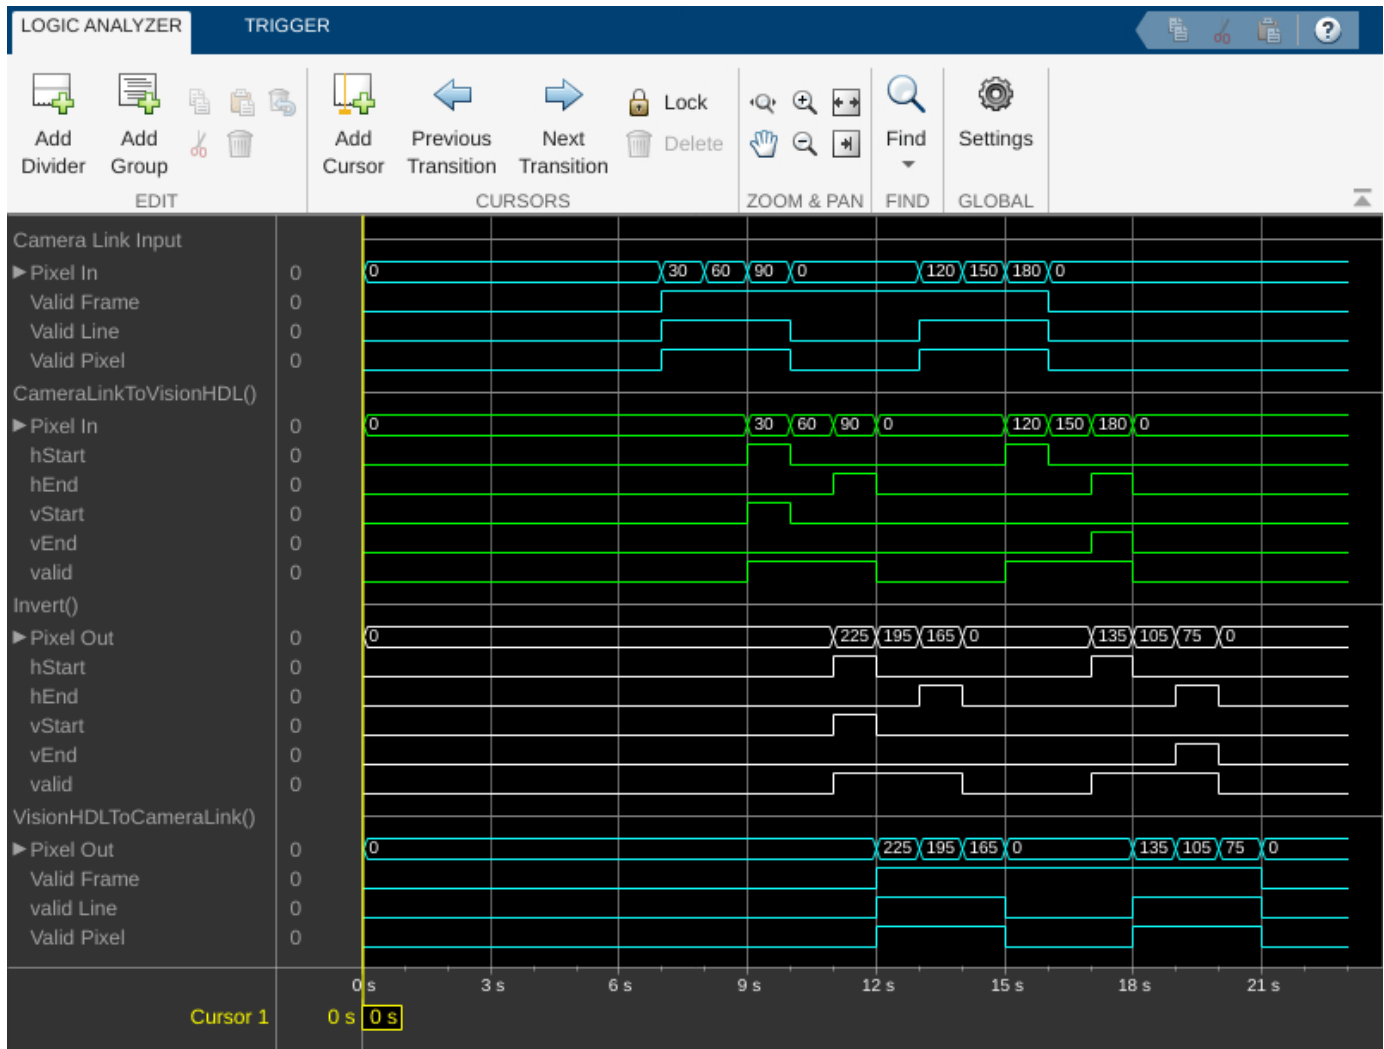
See Also
`pixelcontrolstruct` | `pixelcontrolsignals`

More About
- "Streaming Pixel Interface" on page 1-2

Integrate Vision HDL Blocks into Camera Link System

This example shows how to design a Vision HDL Toolbox™ algorithm for integration into an existing system that uses the Camera Link® signal protocol.

Vision HDL Toolbox blocks use a custom streaming video format. If you integrate Vision HDL Toolbox algorithms into existing design and verification code that operates in a different streaming video format, you must convert the control signals at the boundaries. The example uses custom System objects to convert the control signals between the Camera Link format and the Vision HDL Toolbox `pixelcontrol` format. The model imports the System objects to Simulink® by using the MATLAB® System block.

Structure of the Model

This model imports pixel data and control signals in the Camera Link format from the MATLAB workspace. The `CameraLink_InvertImage` subsystem is designed for integration into existing systems that use Camera Link protocol. The `CameraLink_InvertImage` subsystem converts the control signals from the Camera Link format to the `pixelcontrol` format, modifies the pixel data using the Lookup Table block, and then converts the control signals back to the Camera Link format. The model exports the resulting data and control signals to workspace variables.
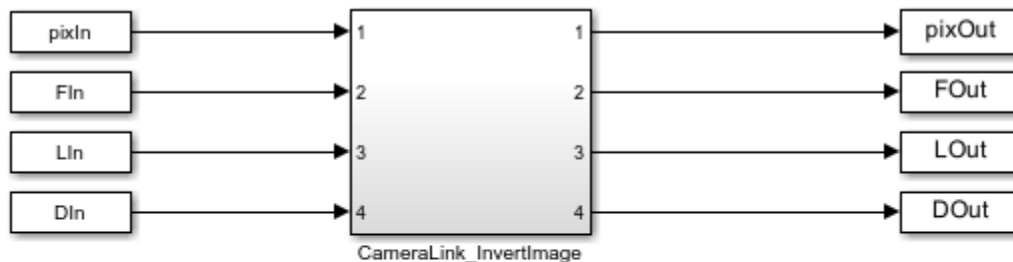
Copyright 2016-2022 The MathWorks, Inc.

Structure of the Subsystem

The `CameraLink2VHT` and `VHT2CameraLink` blocks are MATLAB System blocks that point to custom System objects. The objects convert between Camera Link signals and the `pixelcontrol` format used by Vision HDL Toolbox blocks and objects.

You can put any combination of Vision HDL Toolbox blocks into the middle of the subsystem. This example uses an inversion Lookup Table.

You can generate HDL from this subsystem.
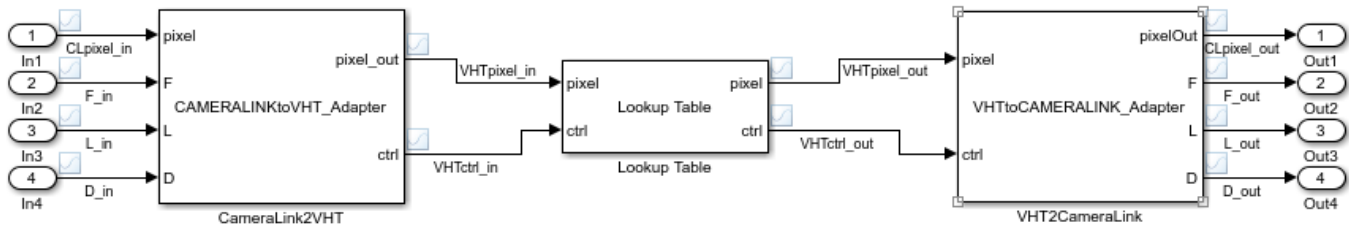
blocks do not need to know the size/format of the frame

Import Data in Camera Link Format

Camera Link consists of three control signals: F indicates the valid frame, L indicates each valid line, and D indicates each valid pixel. For this example, the input data and control signals are defined in the `InitFcn` callback. The vectors describe this 2-by-3, 8-bit grayscale frame. In the figure, the active image area is in the dashed rectangle, and the inactive pixels surround it. The pixels are labeled with their grayscale values.
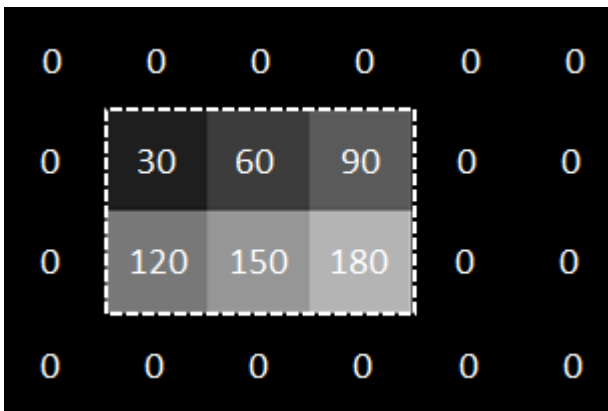

```
FIn = logical([0,0,0,0,0,0,0,1,1,1,1,1,1,1,1,1,1,0,0,0,0,0,0,0,0]);
LIn = logical([0,0,0,0,0,0,0,1,1,1,0,0,0,1,1,1,0,0,0,0,0,0,0,0,0]);
DIn = logical([0,0,0,0,0,0,0,1,1,1,0,0,0,1,1,1,0,0,0,0,0,0,0,0,0]);
pixIn = uint8([0,0,0,0,0,0,0,30,60,90,0,0,0,120,150,180,0,0,0,0,0,0,0,0,0]);
```

Convert Camera Link Control Signals to pixelcontrol Format

Write a custom System object to convert Camera Link signals to the Vision HDL Toolbox format. This example uses the object designed in the "Convert Camera Control Signals to pixelcontrol Format" on page 1-26 example.

The object converts the control signals, and then creates a structure that contains the new control signals. When the object is included in a MATLAB System block, the block translates this structure into the bus format expected by Vision HDL Toolbox blocks. For the complete code for the System object, see `CAMERALINKtoVHT_Adapter.m`.

Create a MATLAB System block and point it to the System object.

Block Parameters: MATLAB System

MATLAB System

Implement block using a System object. Specify the class name.

System object name: CAMERALINKtoVHT_Adapter

New

OK   Cancel   Help
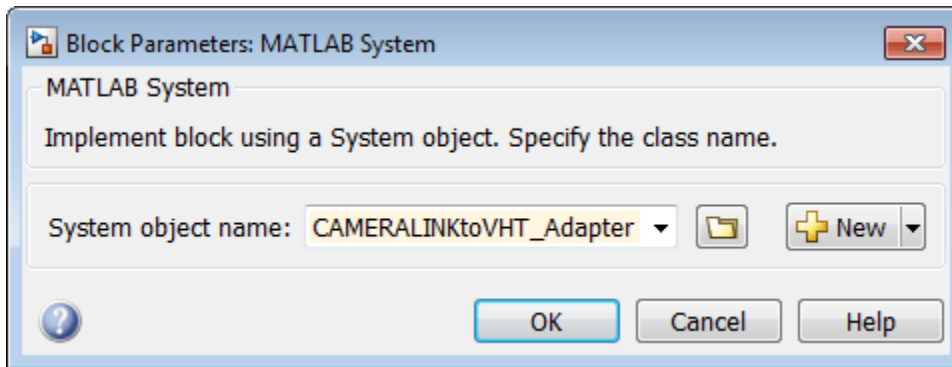
Design Vision HDL Toolbox Algorithm

Select Vision HDL Toolbox blocks to process the video stream. These blocks accept and return a scalar pixel value and a `pixelcontrol` bus that contains the associated control signals. This standard interface makes it easy to connect blocks from the Vision HDL Toolbox libraries together.

This example uses the Lookup Table block to invert each pixel in the test image. Set the table data to the reverse of the `uint8` grayscale color space.
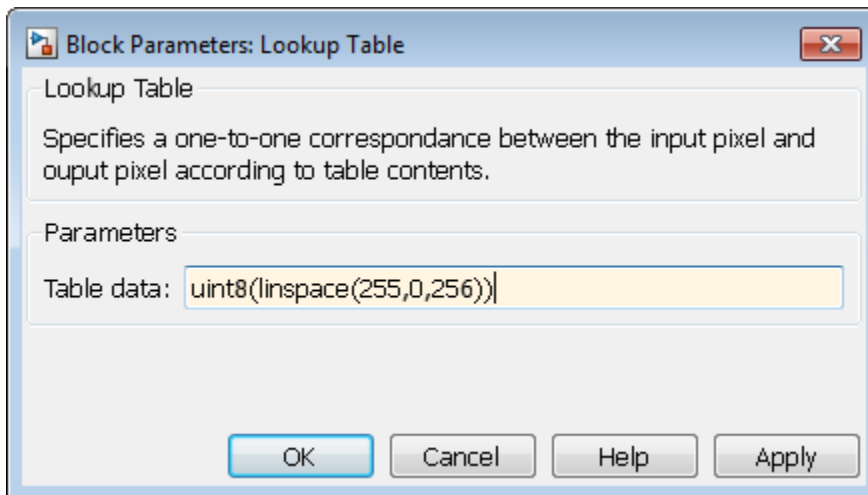
Block Parameters: Lookup Table

Lookup Table

Specifies a one-to-one correspondance between the input pixel and ouput pixel according to table contents.

Parameters

Table data: uint8(linspace(255,0,256))

OK   Cancel   Help   Apply

Convert pixelcontrol to Camera Link

Write a custom System object to convert Vision HDL Toolbox signals back to the Camera Link format. This example uses the object designed in the "Convert Camera Control Signals to pixelcontrol Format" on page 1-26 example.

The object accepts a structure of control signals. When you include the object in a MATLAB System block, the block translates the input `pixelcontrol` bus into this structure. Then it computes the equivalent Camera Link signals. For the complete code for the System object, see `VHTtoCAMERALINKAdapter.m`.

Create a second MATLAB System block and point it to the System object.

View Results

Run the simulation. The resulting vectors represent this inverted 2-by-3, 8-bit grayscale frame. In the figure, the active image area is in the dashed rectangle, and the inactive pixels surround it. The pixels are labeled with their grayscale values.
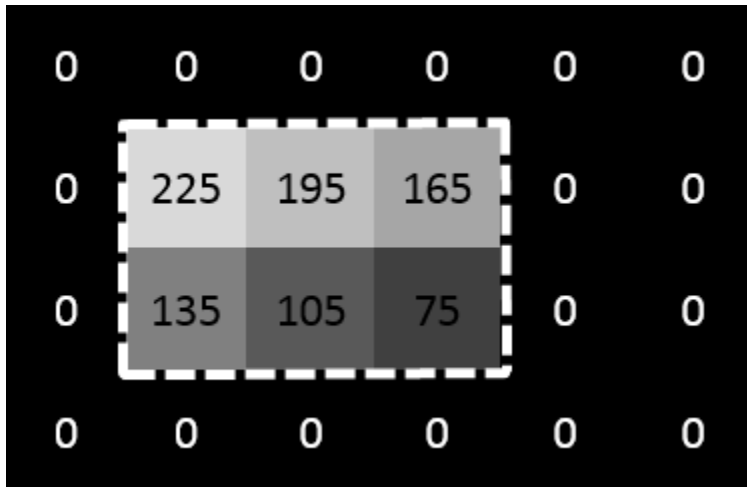
If you have a DSP System Toolbox™ license, you can view the signals over time using the Logic Analyzer. Select all the signals in the CameraLink_InvertImage subsystem for streaming, and open the Logic Analyzer. This waveform shows the input and output Camera Link control signals and pixel values at the top, and the input and output of the Lookup Table block in pixelcontrol format at the bottom. The pixelcontrol busses are expanded to observe the boolean control signals.
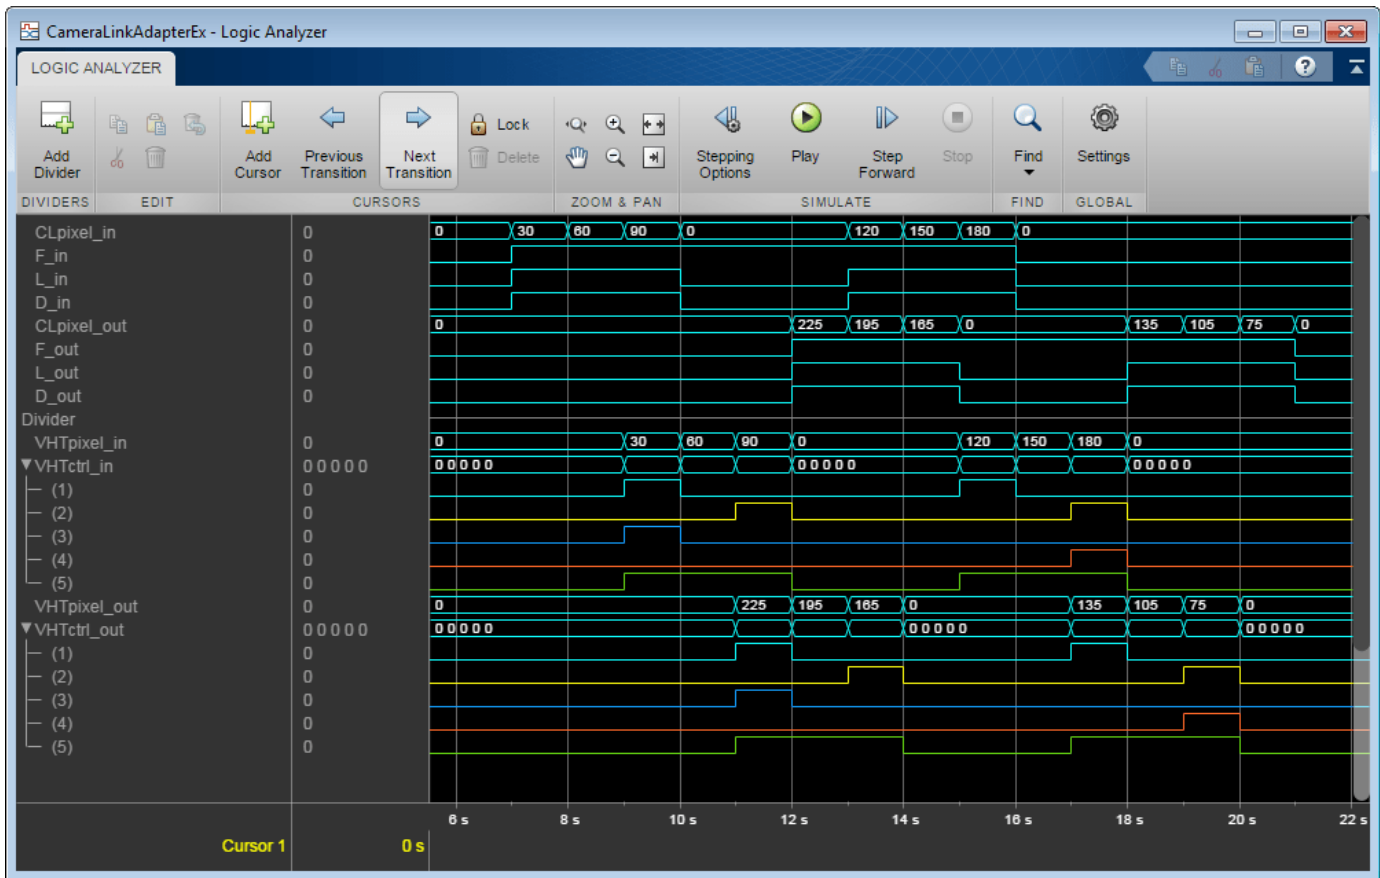
For more info on observing waveforms in Simulink, see "Inspect and Measure Transitions Using the Logic Analyzer" (DSP System Toolbox).

Generate HDL Code for Subsystem

To generate HDL code you must have an HDL Coder™ license.

To generate the HDL code, use the following command.

```
makehdl('CameraLinkAdapterEx/CameraLink_InvertImage')
```

You can now simulate and synthesize these HDL files along with your existing Camera Link system.

See Also

More About

• "Streaming Pixel Interface" on page 1-2

HDL-Optimized Algorithm Design

Configure Blanking Intervals
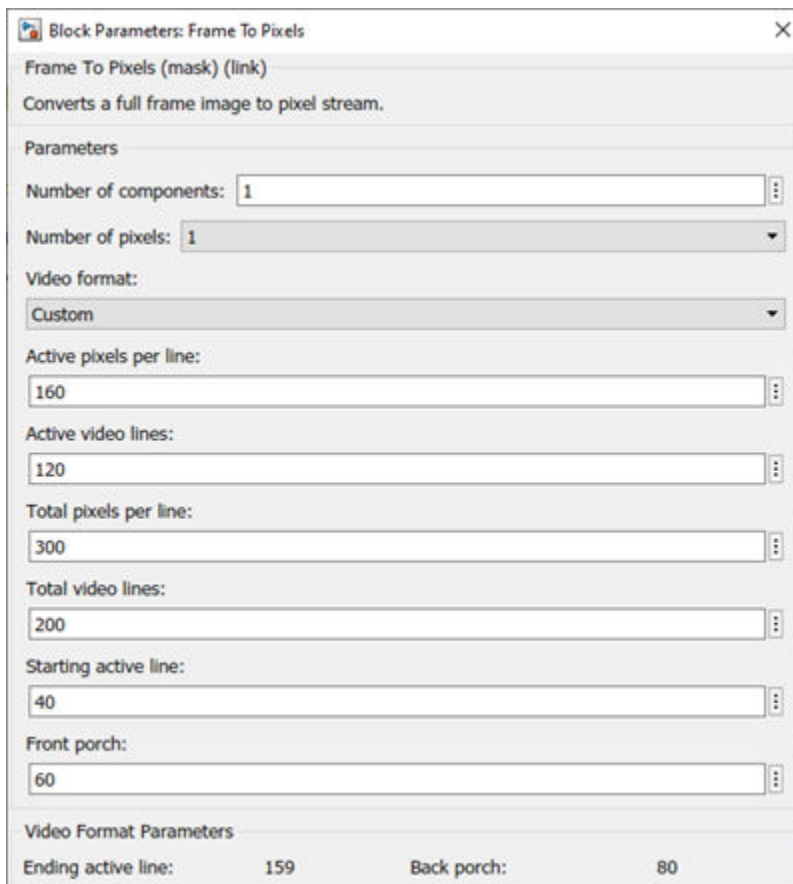
Streaming video protocols have two blanking intervals: horizontal and vertical. The horizontal blanking interval is the period of inactive cycles between the end of one line and the beginning of the next line. The vertical blanking interval is the period of inactive lines between the end of a frame and the beginning of the next frame.

In this frame diagram, the blue shaded areas to the left and right of the active frame indicate the horizontal blanking interval. The orange shaded areas above and below the active frame indicate the vertical blanking interval.
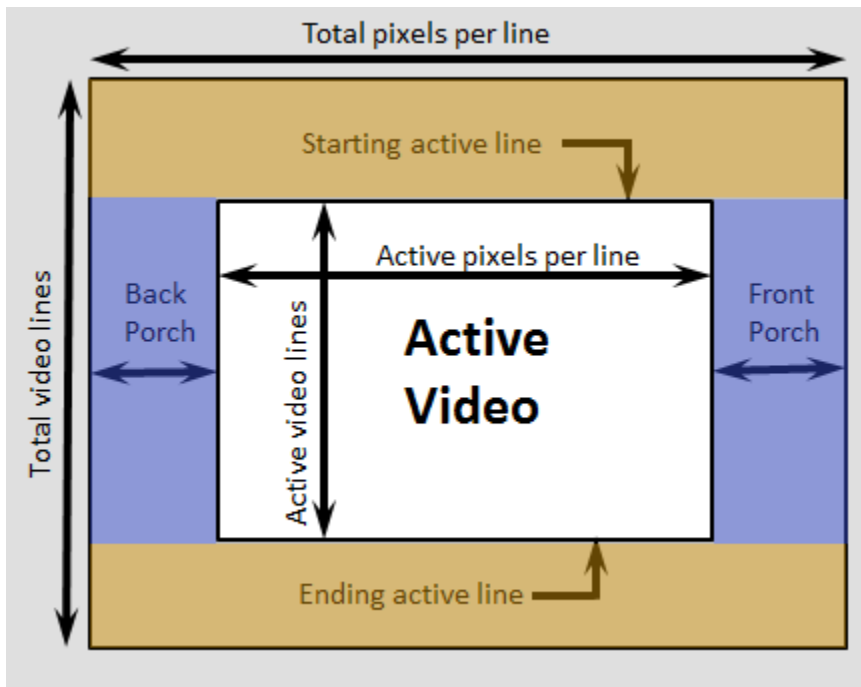
In the Frame To Pixels block, the horizontal blanking interval is equal to **Total pixels per line** – **Active pixels per line** or, equivalently, **Front porch** + **Back porch**. The vertical blanking interval is equal to **Total video lines** – **Active video lines** or, equivalently, **Starting active line** + **Ending active line** – **Active video lines**.

For example, the Frame To Pixels block whose parameters are shown in this image has a horizontal blanking interval of 140 pixels and a vertical blanking interval of 80 lines.

Block Parameters: Frame To Pixels

Frame To Pixels (mask) (link)

Converts a full frame image to pixel stream.

Parameters

Number of components: 1

Number of pixels: 1

Video format:

Custom

Active pixels per line:

160

Active video lines:

120

Total pixels per line:

300

Total video lines:

200

Starting active line:

40

Front porch:

60

Video Format Parameters

Ending active line: 159   Back porch: 80

A streaming video format must have a long enough blanking interval so that the operation on the previous line or frame completes before the next line or frame starts. An inadequate horizontal or vertical blanking interval results in corrupted output frames. Standard streaming video formats use a horizontal blanking interval of about 25% of the line width. This interval is much larger than the delay of a typical operation. However, when you use a custom video format, you must include blanking intervals that accommodate the length of the operations in your design.

In these waveform diagrams, the top signal shows the state of the pixel stream for two lines of a frame. The shaded area represents the horizontal blanking interval between lines. The bottom signal shows the state of the block performing an operation on the pixel stream. The busy state indicates when the block is processing a line, and the idle state indicates when the block is available to start working on a new line. The first pair of signals shows a scenario where the block finishes working on the first active line before the second line begins. This blanking interval is long enough to ensure correct output frames, because the block is available to start work on the second line when it arrives. The second pair of signals shows a scenario where the block is still working on the first active line when the second line begins. The output of the block is corrupted in the second case, because the block misses the beginning of the second line.
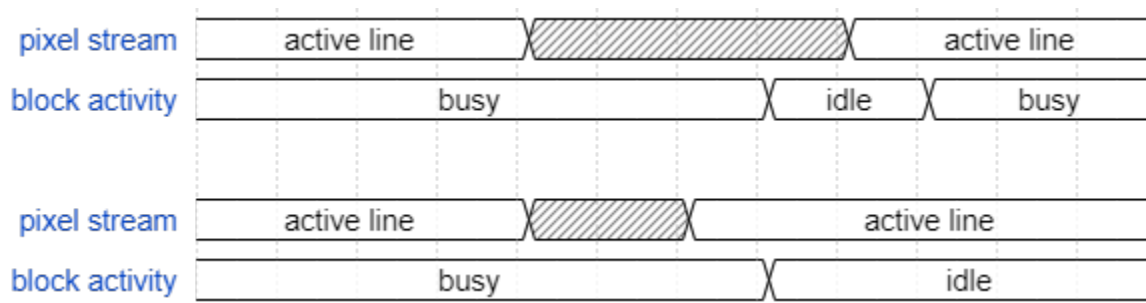
The time an operation takes to complete after the end of the line is often dependent on the kernel size of the operation. For instance, algorithms that use line buffers and apply padding pixels to the edge of the frame require at least *Kw* cycles between lines, where *Kw* is the width of the kernel. An algorithm might also have pipeline delays from the kernel operation after the buffer. These delays can be related or unrelated to the kernel size, and can be greater or smaller than the line buffer delays. The processing time of each operation depends on the line buffer pipelining and on the kernel operation pipelining. The blanking interval must be long enough to accommodate the longer of these two delays. When you use multiple blocks in a processing chain, the blanking interval must accommodate the block with the longest delay.

The recommended minimum horizontal blanking interval is 2×*Kw* when using padding or 12 cycles when you set the **Padding method** parameter to `None`. This interval includes some margin for longer kernel processing times on top of the line buffer delay.

The recommended vertical blanking interval is at least the height of the kernel, *Kh* lines. The line buffer requires this interval whether or not the operation uses padding.

Note   When you use a pixel streaming block inside an Enabled Subsystem, the enable signal pattern must maintain the timing of the pixel stream, including the minimum blanking intervals. You may need to extend the blanking intervals to accommodate for cycles when the enable is low.

Troubleshoot Blanking Interval Problems

When the blanking interval is too small, you might see:

- Blank output frames
- Partial output frames
- Corrupted pixel stream control signal patterns (for instance, missing `vEnd` or `hEnd` signals, or duplicate `End` or `Start` signals)
- That the algorithm works with continuous valid input pixels on each line, but not when gaps exist between valid pixels in a line
- That the algorithm works in Simulink but fails in HDL simulation

Vision HDL Toolbox library blocks model hardware pipeline stages as a latency applied at the output. In the corresponding HDL implementations, the pipeline stages are distributed across the calculation. This difference means that for a given cycle, a block can be in a busy state in HDL simulation but appear idle in Simulink. When the blanking period is too short, this difference can cause the generated HDL test bench to show mismatches between Simulink and HDL signals, especially on the output control signals.

If you see any of these symptoms, increase your horizontal and vertical blanking intervals to 25% of the active frame dimensions and rerun the simulation. If this step confirms that a too-small blanking interval is causing your symptoms, you can fine tune the intervals.

One way to diagnose blanking interval problems in Simulink is to use the Measure Timing block to observe the dimensions of the pixel stream before and after your operation. Inadequate blanking intervals cause the block to corrupt the control signals. In these cases, the output frames show different dimensions than the input frames.

This model shows an Image Filter block configured with a 12-by-12 filter kernel and edge padding enabled. The pixel stream format is a custom format that has only 8 horizontal blanking pixels, as shown by the Measure Timing block on the input stream. Because the horizontal blanking interval is smaller than the kernel width, the output frame is blank. The Measure Timing block on the output of the filter shows corruption of the format.
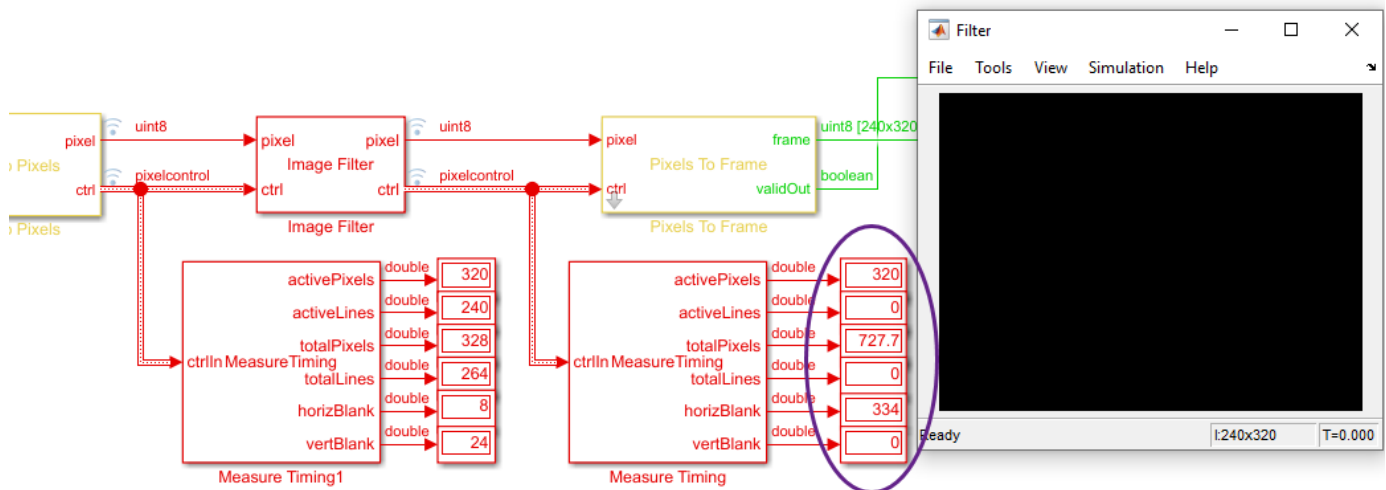
You can also see the corruption by looking at the output control signals in the **Logic Analyzer** app. The waveform shows the input and output signals of the Image Filter block. The red arrows indicate missing `hStart` signals and a different pattern on the output `valid` signal from the block.
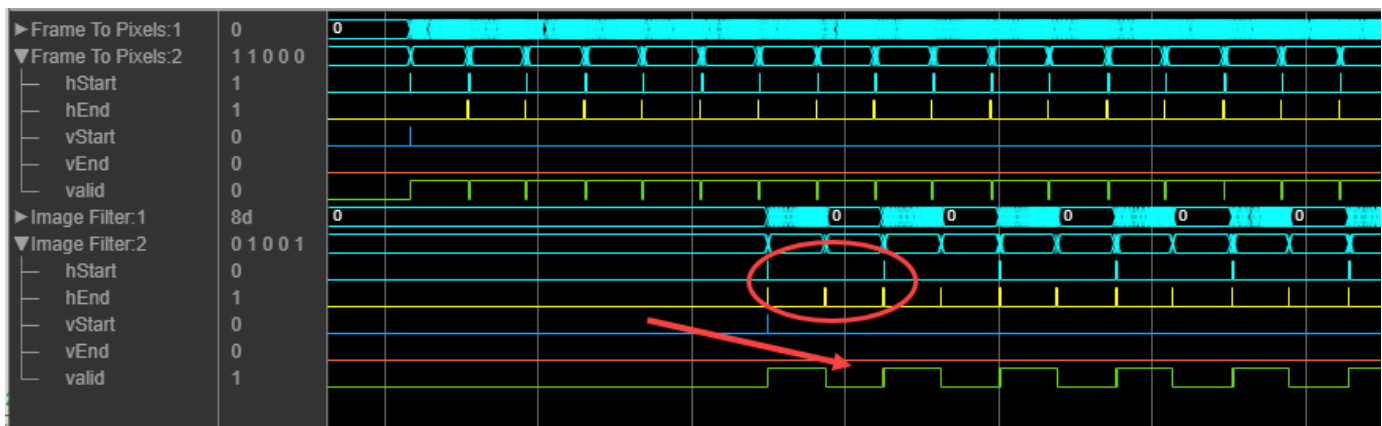
This model shows an Edge Detector block configured to use the 3-by-3 Sobel filter kernel and with edge padding enabled. This pixel stream format has only two horizontal blanking pixels, as shown by

the Measure Timing block on the input stream. In this case, the output frame includes only every second line. The Measure Timing block on the output of the filter shows the corruption of the format.
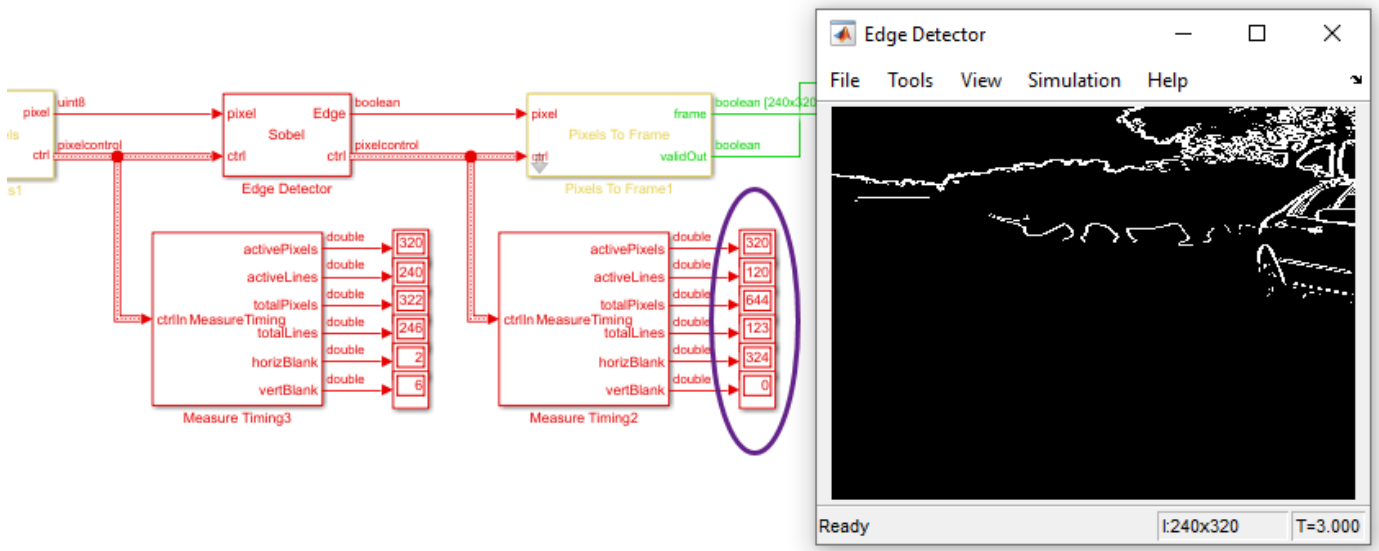
You can also see the corruption by looking at the output control signals in the **Logic Analyzer** app. The waveform shows the input and output signals of the Edge Detector block. The red circle indicates missing `hStart` and `hEnd` signals, and the red arrow indicates a different pattern on the output `valid` signal from the block.
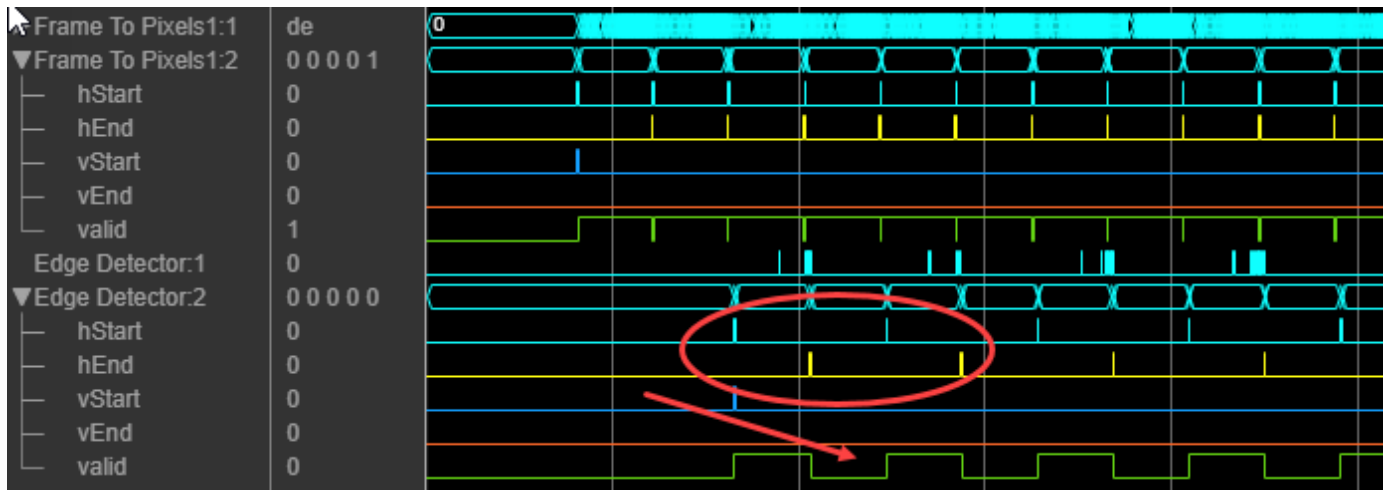
If you modify the input format to have a horizontal blanking interval of 3 pixels, this model returns the correct output frames in Simulink. However, when you run the generated HDL test bench, the test bench reports mismatches between the signals captured in Simulink and the signal behavior in HDL. This image of the test bench log highlights the mismatch in the output `hEnd` and `vEnd` signals.

```
# ERROR in Out2_hStart at time          781920 : Expected '0' Actual '1'
# ERROR in Out1 at time          781920 : Expected '0' Actual '1'
# ERROR in Out2_valid at time          781920 : Expected '0' Actual '1'
# ERROR in Out2_hStart at time          781930 : Expected '1' Actual '0'
# ERROR in Out1 at time          784750 : Expected '1' Actual '0'
# ERROR in Out1 at time          784990 : Expected '0' Actual '1'
# ERROR in Out1 at time          785100 : Expected '1' Actual '0'
# ERROR in Out2_valid at time          785100 : Expected '1' Actual '0'
# ERROR in Out2_valid at time          785110 : Expected '1' Actual '0'
# ERROR in Out1 at time          785110 : Expected '1' Actual '0'
# ERROR in Out2_vEnd at time          785120 : Expected '1' Actual '0'
# ERROR in Out2_hEnd at time          785120 : Expected '1' Actual '0'
# ERROR in Out1 at time          785120 : Expected '1' Actual '0'
# ERROR in Out2_valid at time          785120 : Expected '1' Actual '0'
```

This waveform from the simulation of the HDL test bench shows that the hEnd and vEnd signals at the end of the first frame are missing. The blue signals are the expected output as captured from the Simulink simulation. The red signals are the output of the algorithm in the HDL simulation. The red arrows indicate where the expected control signal pulses are missing.
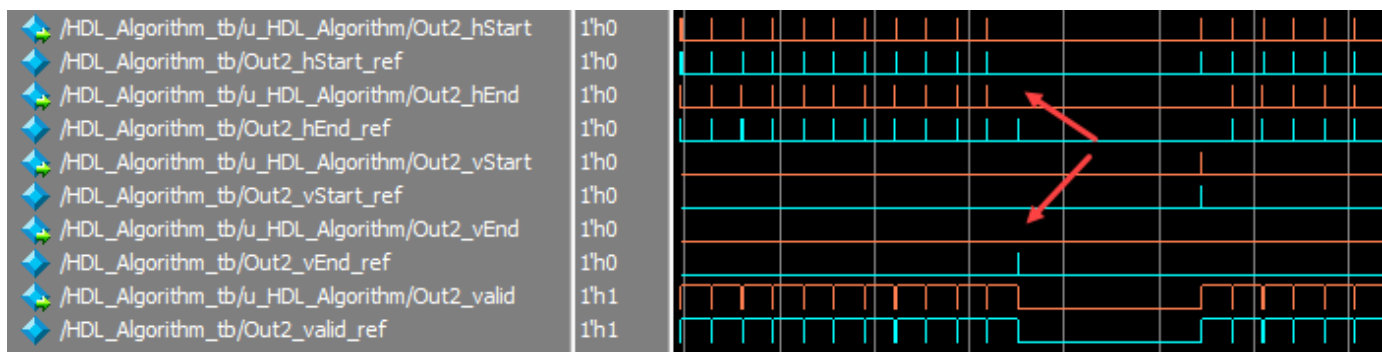
To fix the Image Filter model and the Edge Detector model, set the horizontal blanking interval to at least 2×*Kw* pixels, where *Kw* is the width of the filter kernel. For the Image Filter model, set this value to at least 24 pixels. For the Edge Detector model, set this value to at least 8 pixels.

See Also

Frame To Pixels | Measure Timing

More About

*   "Streaming Pixel Interface" on page 1-2
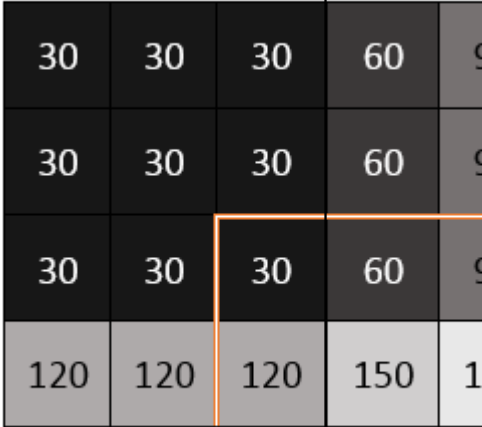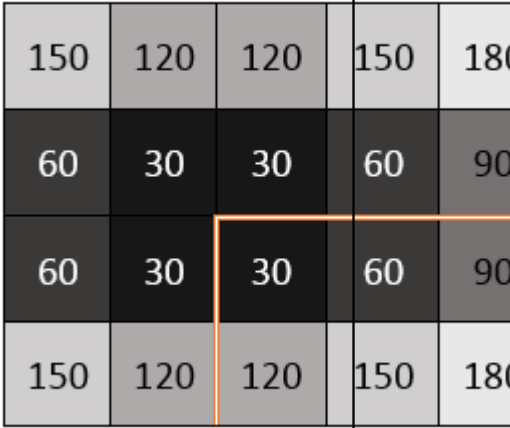
Edge Padding

To perform a kernel-based operation such as filtering on a pixel at the edge of a frame, Vision HDL Toolbox algorithms pad the edges of the frame with extra pixels. These padding pixels are used for internal calculation only. The output frame has the same dimensions as the input frame. The padding operation assigns a pattern of pixel values to the inactive pixels around a frame. Vision HDL Toolbox algorithms provide padding by constant value, replication, or symmetry.

Some blocks and System objects also support opting out of setting the padding pixel values. This option reduces the hardware resources used by the block and the blanking required between frames but affects the accuracy of the output pixels at the edges of the frame.

The diagrams show the top-left corner of a frame, with padding added to accommodate a 5-by-5 filter kernel. When computing the filtered value for the top-left active pixel, the algorithm requires two rows and two columns of padding. The edge of the active image is indicated by the double line.

Type of Padding	Description	Diagram
Constant	Each added pixel is assigned the same value. On some blocks and System objects you can specify the constant value. The value 0, representing black, is a reserved value in some video standards. Choosing a small number, such as 16, as a near-black padding value, is common.	In the diagram, C is the constant value assigned to the inactive pixels around the active frame.

Type of Padding	Description	Diagram
Replicate	The pixel values at the edge of the active frame are repeated to make rows and columns of padding pixels.	The diagram shows the pattern of replicated values assigned to the inactive pixels around the active frame.
Symmetric	The padding pixels are added such that they mirror the edge of the image.	The diagram shows the pattern of symmetric values assigned to the inactive pixels around the active frame. The pixel values are symmetric about the edge of the image in both dimensions.

Type of Padding	Description	Diagram
Reflection	The padding pixels are added such that they reflect around the pixel at the edge of the image. This type of padding is useful for machine learning applications because it removes edge contrast and maintains texture.	The diagram shows the pattern of reflected values assigned to the inactive pixels around the active frame. The pixel values are reflected around the first pixel of the image in both dimensions. 330 300 270 300 330 210 180 150 180 210 90 60 30 60 90 210 180 150 180 210 330 300 270 300 330
None	This option excludes padding logic. The line buffer does not set the pixels outside the image frame to any particular value. The kernel calculation uses the current value in the line buffer. To maintain pixel stream timing, the output frame is the same size as the input frame. However, to avoid using pixels calculated from undefined padding values, mask off the *KernelSize*/2 pixels around the edge of the frame for downstream operations. Excluding padding can useful for applications that meet any of these conditions. • The output video stream does not need to maintain physical timing. • The resulting image is not displayed. For example, finding the location of objects in an image. • The information of interest is always in the center of the image. For an example, see "Increase Throughput by Omiting Padding" on page 2-12.	The diagram shows the undefined values of the inactive pixels around the active frame. ? ? ? ? ? ? ? ? ? ? ? ? 30 60 90 ? ? 120 150 180

Padding requires minimum horizontal and vertical blanking periods. This interval gives the algorithm time to add and store the extra pixels. The blanking period, or inactive pixel region, must be at least *KernelWidth* pixels horizontally and *KernelHeight* lines vertically.

When you set the **Padding method** to `None`, the horizontal blanking period must have at least 6 pixels of front porch and 6 pixels of back porch. For the Median Filter block with the **Padding method** set to `None`, the horizontal blanking must have at least 10 pixels of front porch and 10 pixels of back porch. The vertical blanking still must be *KernelHeight* lines. For more detail on blanking intervals, see "Configure Blanking Intervals" on page 2-2.

See Also
Image Filter | `visionhdl.ImageFilter`

More About
- "Streaming Pixel Interface" on page 1-2

Increase Throughput by Omiting Padding

This example shows how to reduce latency and save hardware resources by not adding padding pixels at the edge of each frame.

Most image filtering operations pad the image to fill in the neighborhoods for pixels at the edge of the image. Padding can help avoid border artifacts in the output image. In a hardware implementation, the padding operation uses extra resources and introduces extra latency.

Vision HDL Toolbox™ blocks that perform neighborhood processing with padding require horizontal blanking that is twice the kernel width. This behavior means that larger filter sizes result in a longer blanking requirement. Excluding the padding by setting the **Padding method** parameter to `None` enables you to use a smaller period of horizontal blanking. Without padding, the horizontal blanking requirement is independent of the image resolution and kernel size. A small number of blanking cycles are still required.

This example includes two models. The first model shows how to use this option with library blocks, and the second model demonstrates using it when constructing algorithms that use the Line Buffer block. This example also explains some design considerations when you do not use padding.

Omitting Padding with Library Blocks

This example model shows how to omit padding with a predefined algorithm from Vision HDL Toolbox libraries. This model includes an Image Filter block configured for an n-by-n blur filter and with its **Padding method** parameter set to `None`. You can change the size of the filter kernel by changing the value of `n` in the workspace. The model opens with `n` set to 15.

When using edge padding, most blocks have `floor(KernelHeight/2)` lines of latency and require `2*KernelWidth` cycles of horizontal blanking. When you omit padding, most blocks require only 12 cycles of horizontal blanking. Because the internal line buffer latency no longer depends on the kernel size, this blanking interval accommodates any kernel size.

To show the reduced blanking requirements of using **Padding method** set to `None`, the Frame To Pixels block is configured for a custom 240p format that uses only 12 cycles of combined front and back porch.

When you run the model, it shows these three figures.

- Input Video -- Original 240p input video.
- Padding None Full Frame -- Output video from the filter without padding, showing border artifacts.
- Padding None ROI -- Output video from the filter without padding, with border pixels trimmed from the edges of the frame. The frame size is smaller than the size of the input video.
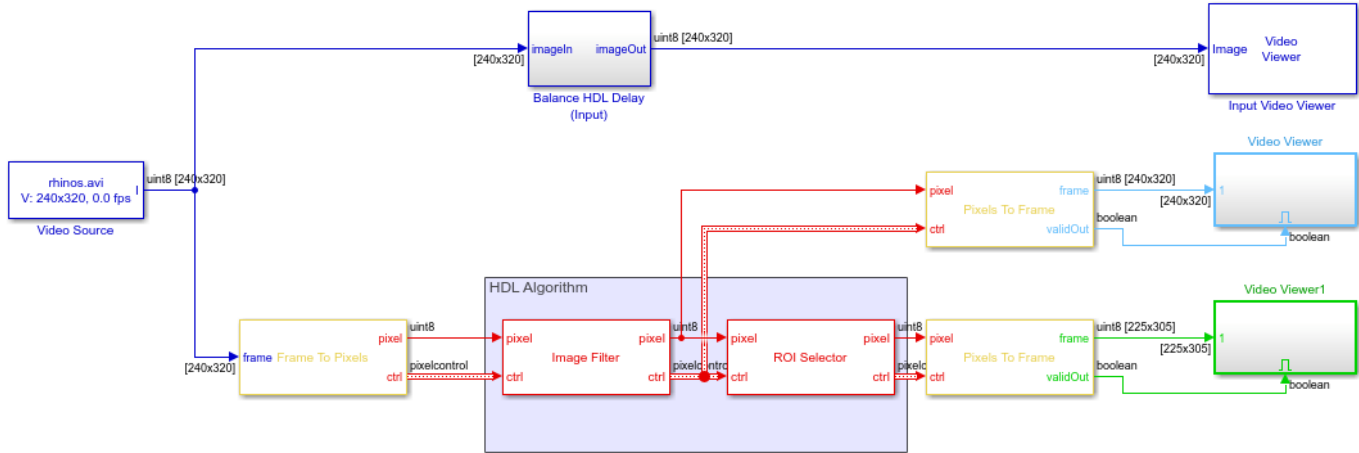
Copyright 2019 The MathWorks, Inc.

Border Artifacts

In the Padding None Full Frame viewer, shown, a dark border is visible around the edge of each frame. This effect is because, without padding pixels, the filter neighborhoods are not fully defined at the edges of the frame. Output from a filter that has padding pixels does not show any border artifacts because the padding logic ensures that the edge neighborhoods are fully defined.

Removing or masking off these border pixels from nonpadded output before further analysis is common. Border artifacts can decrease the accuracy of subsequent processing. For example, these artifacts can affect the statistical distribution of the overall image. Vision HDL Toolbox blocks return the border pixels for nonpadded images to maintain the input and output timing. The values of these pixels are undefined and cannot be assumed to have any particular relation to the surrounding pixels.

The ROI Selector block removes `floor(KernelHeight/2)` and `floor(KernelWidth/2)` pixels from the edges of each frame. The Padding None ROI viewer, shown, shows the video with the border artifacts removed. The resulting frame for a 15-by-15 kernel is 225-by-305 pixels in size, reduced from 240-by-320 pixels.

Omitting Padding with the Line Buffer Block

This model shows how to design algorithms by using a Line Buffer with the **Padding method** parameter set to None. This model contains a Padding None subsystem, and a Padding Symmetric subsystem.

The Frame To Pixels block connected to the Padding Symmetric subsystem uses the standard 240p format. The standard horizontal blanking (combined front and back porch) is 82 cycles. Increasing the resolution increases the blanking interval. For example, the 1080p format has 280 idle cycles between lines.

The Frame To Pixels block connected to the Padding None subsystem implements a custom 240p format that uses only 12 cycles of combined front and back porch, the same as in the Image Filter model shown earlier.

This model implements a 15-by-15 Gaussian filter, with a large standard deviation, by using the Line Buffer block.

When you run the model, it shows three figures:

- Input Video -- Original 240p input video.
- Padding None ROI -- Output video from the filter without padding, with border pixels trimmed from the edges of the frame. The frame size is smaller than the size of the input video.
- Padding Symmetric -- Output video from the filter with symmetric padding. This video is full size but has no edge effects because the padding bits define the neighborhoods around the edge pixels.
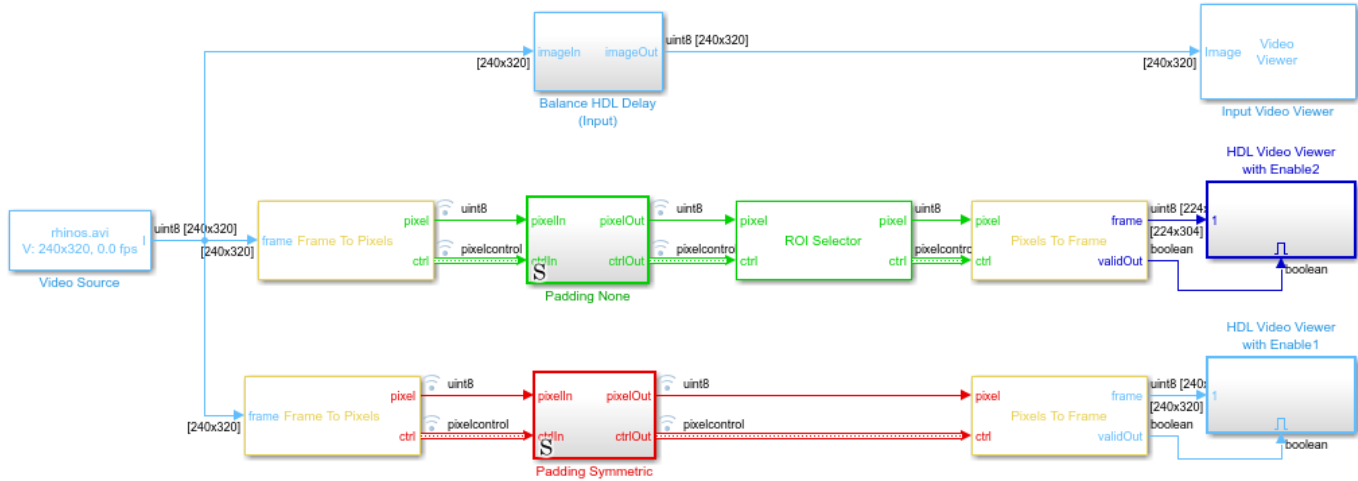
Copyright 2019 The MathWorks, Inc.

`pixelcontrol` Delay Balancing

When you construct algorithms that use the Line Buffer block, you must delay-balance the `pixelcontrol` bus to account for the kernel latency. When you use padding, the Line Buffer returns **shiftEnable** set to 1 for `floor(KernelWidth/2)` cycles before **hStart** and after **hEnd**. The delay-balancing logic uses this extended **shiftEnable** signal to control the delay registers for the `pixelcontrol` signals. You can see this logic in the Padding Symmetric/pixelctrldelay subsystem.

When you set **Padding method** to None, the Line Buffer returns **shiftEnable** to 1 between **hStart** and **hEnd**. The delay-balancing logic must use the clock, instead of **shiftEnable**, to control the delay registers for **hEnd**, **vEnd**, and **valid**. The **valid** signal must also respond to **shiftEnable** being set to 0 during a line, which can occur when interfacing with external memory. The **valid** signal must also be set to 1 on the last pixel of the line, to match with **hEnd** and **vEnd**. To meet both requirements, the delay-balancing logic delays the **valid** signal by using a register enabled by **shiftEnable**, and uses a Unit Delay Enabled block to set the **valid** signal to 1 with **hEnd** at the end of the line. The Padding None/pixelctrldelay subsystem shows this logic.
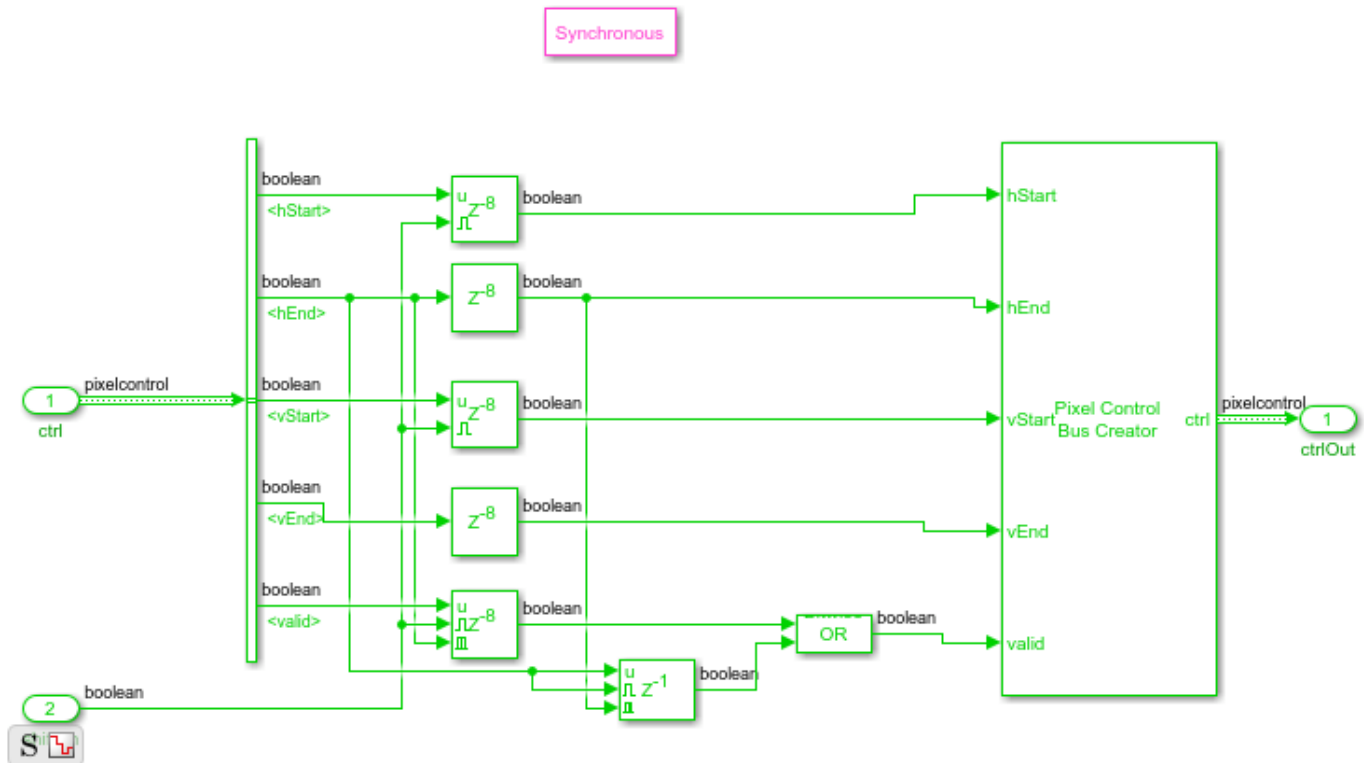
Conclusion

Excluding padding logic enables you to achieve higher throughput by using a video format with reduced horizontal blanking. This option also reduces hardware resource usage. However, your design must account for the border artifacts later in the processing chain. When you use the Line Buffer block, you must delay the `pixelcontrol` bus to match the kernel latency by using control logic that accounts for the modified behavior of the **shiftEnable** output signal. Using this example as a starting point, you can design algorithms and systems that achieve higher throughput by excluding padding logic.

Gamma Correction

This example shows how to model pixel-streaming gamma correction for hardware designs. The model compares the results from the Vision HDL Toolbox™ Gamma Corrector block with the results generated by the full-frame Gamma block from Computer Vision Toolbox™.

This example model provides a hardware-compatible algorithm. You can implement this algorithm on a board using a Xilinx™ Zynq™ reference design. See "Gamma Correction with Zynq-Based Hardware" (Vision HDL Toolbox Support Package for Xilinx Zynq-Based Hardware).

Structure of the Example

The Computer Vision Toolbox product models at a high level of abstraction. The blocks and objects perform full-frame processing, operating on one image frame at a time. However, FPGA or ASIC systems perform pixel-stream processing, operating on one image pixel at a time. This example simulates full-frame and pixel-streaming algorithms in the same model.

The GammaCorrectionHDL.slx system is shown below.
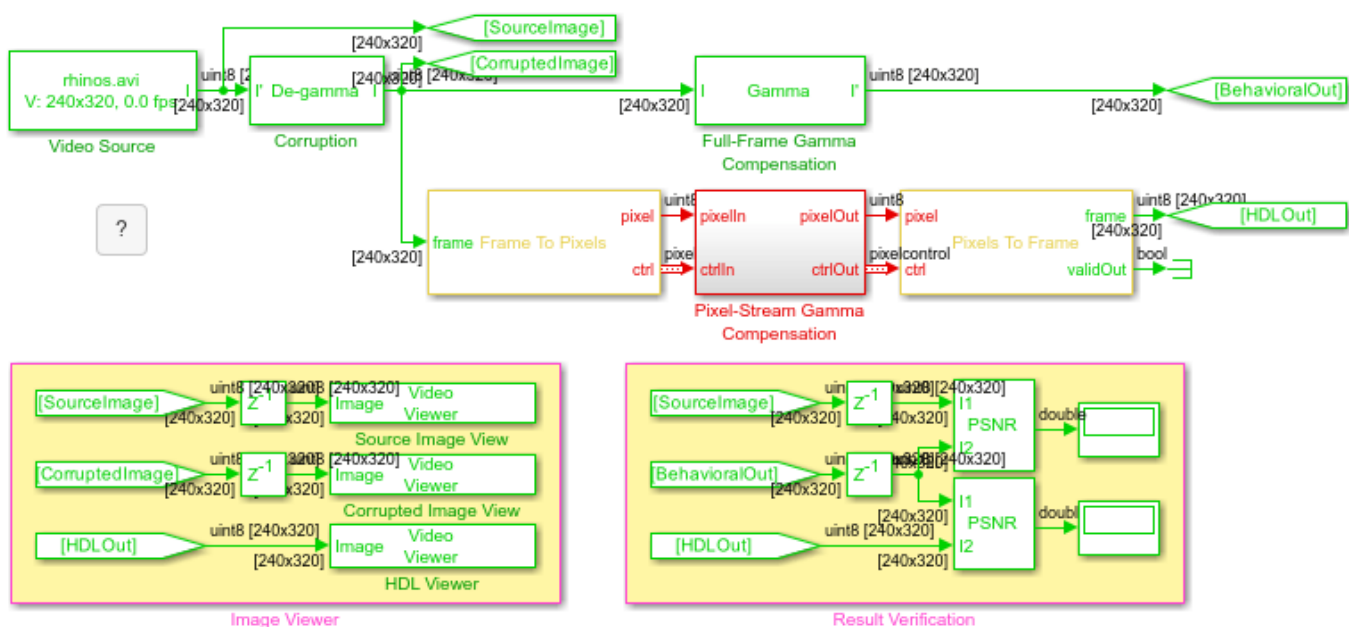
Copyright 2018 The MathWorks, Inc.

The difference in the color of the lines feeding the **Full-Frame Gamma Compensation** and **Pixel-Stream Gamma Compensation** subsystems indicates the change in the image rate on the streaming branch of the model. This rate transition is because the pixel stream is sent out in the same amount of time as the full video frames and therefore it is transmitted at a higher rate.

In this example, the Gamma correction is used to correct dark images. Darker images are generated by feeding the **Video Source** to the **Corruption** block. The **Video Source** outputs a 240p grayscale video, and the **Corruption** block applies a De-gamma operation to make the source video perceptually darker. Then, the downstream **Full-Frame Gamma Compensation** block or **Pixel-**

Stream Gamma Compensation subsystem removes the previous De-gamma operation from the corrupted video to recover the source video.

One frame of the source video, its corrupted version, and recovered version, are shown from left to right in the diagram below.
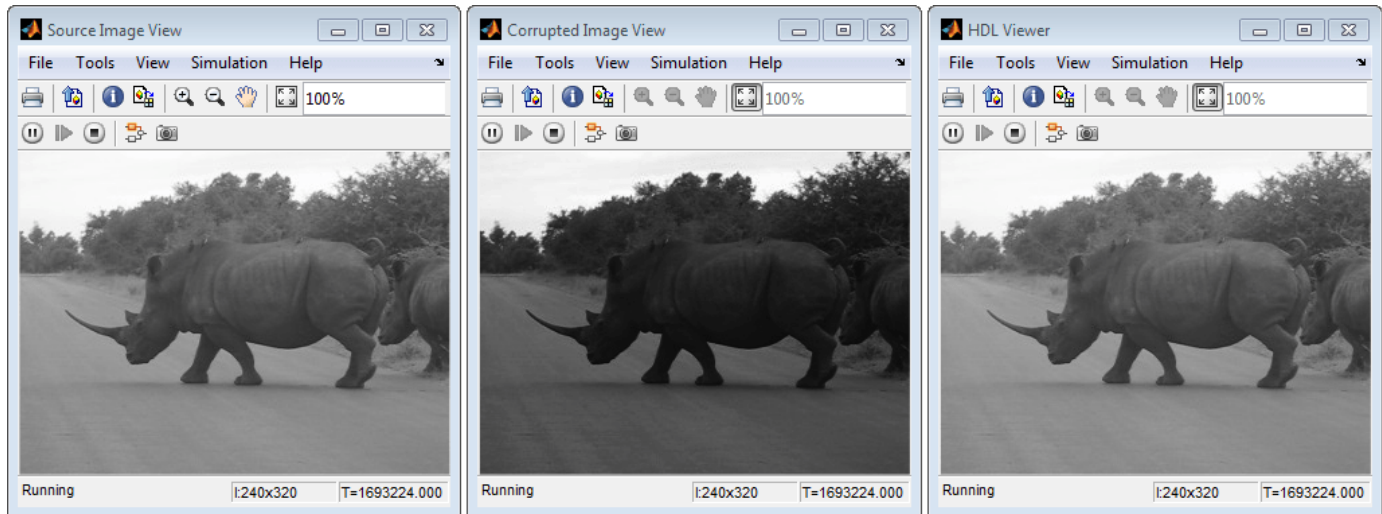
It is a good practice to develop a behavioral system using blocks that process full image frames, the **Full-Frame Gamma Compensation** block in this example, before moving forward to working on an FPGA-targeting design. Such a behavioral model helps verify the video processing design. Later on, it can serve as a reference for verifying the implementation of the algorithm targeted to an FPGA. Specifically, the lower **PSNR** (peak signal-to-noise ratio) block in the **Result Verification** section at the top level of the model compares the results from full-frame processing with those from pixel-stream processing.

Frame To Pixels: Generating a Pixel Stream

The task of the **Frame To Pixels** is to convert a full-frame image to pixel stream. To simulate the effect of horizontal and vertical blanking periods found in real life hardware video systems, the active image is augmented with non-image data. For more information on the streaming pixel protocol, see "Streaming Pixel Interface" on page 1-2. The **Frame To Pixels** block is configured as shown:

The **Number of components** field is set to 1 for grayscale image input, and the **Video format** field is 240p to match that of the video source.

In this example, the Active Video region corresponds to the 240x320 matrix of the dark image from the upstream **Corruption** block. Six other parameters, namely, **Total pixels per line**, **Total video**

lines, **Starting active line**, **Ending active line**, **Front porch**, and **Back porch** specify how many non-image data will be augmented on the four sides of the Active Video. For more information, see the Frame To Pixels block reference page.

Note that the sample time of the **Video Source** is determined by the product of **Total pixels per line** and **Total video lines**.

Gamma Correction

As shown in the diagram below, the **Pixel-Stream Gamma Compensation** subsystem contains only a **Gamma Corrector** block.
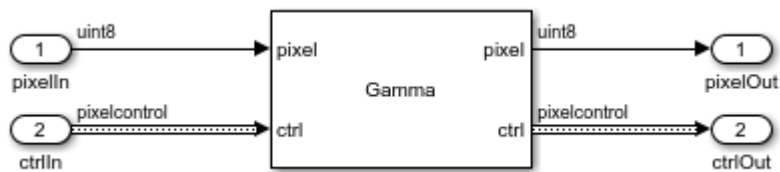
The **Gamma Corrector** block accepts the pixel stream, as well as a bus containing five synchronization signals, from the **Frame To Pixels** block. It passes the same set of signals to the downstream **Pixels To Frame** block. Such signal bundle and maintenance are necessary for pixel-stream processing.

Pixels To Frame: Converting Pixel Stream Back to Full Frame

As a companion to **Frame To Pixels** that converts a full image frame to pixel stream, the **Pixels To Frame** block, reversely, converts the pixel stream back to the full frame by making use of the synchronization signals. Since the output of the **Pixels To Frame** block is a 2-D matrix of a full image, there is no need to further carry on the bus containing five synchronization signals.

The **Number of components** field and the **Video format** fields of both Frame To Pixels and Pixels To Frame are set at 1 and 240p, respectively, to match the format of the video source.

Image Viewer and Result Verification

When you run the simulation, three images will be displayed (refer to the images shown in the "Structure of the Example" Section):

- The source image given by the **Image Source** subsystem
- The dark image produced by the **Corruption** block
- The HDL output generated by the **Pixel-Stream gamma Compensation** subsystem

The presence of the four **Unit Delay** blocks on top level of the model is to time-align the 2-D matrices for a fair comparison.

While building the streaming portion of the design, the **PSNR** block continuously verifies the **HDLOut** results against the original full-frame design **BehavioralOut**. During the course of the simulation, this **PSNR** block should give **inf** output, indicating that the output image from the **Full-Frame Gamma Compensation** matches the image generated from the stream processing **Pixel-Stream Gamma Compensation** model.

Exploring the Example

The example allows you to experiment with different Gamma values to examine their effect on the Gamma and De-gamma operation. Specifically, a workspace variable $gammaValue$ with an initial value 2.2 is created upon opening the model. You can modify its value using the MATLAB command line as follows:

```
gammaValue=4
```

The updated $gammaValue$ will be propagated to the **Gamma** field of the **Corruption** block, the **Full-Frame Gamma Compensation** block, and the **Gamma Corrector** block inside **Pixel-Stream Gamma Compensation** subsystem. Closing the model clears $gammaValue$ from your workspace.

Although Gamma operation is conceptually the inverse of De-gamma, feeding an image to Gamma followed by a De-gamma (or De-gamma first then Gamma) does not necessarily perfectly restore the original image. Distortions are expected. To measure this, in our example, another **PSNR** block is placed between the **SourceImage** and **BehavioralOut**. The higher the PSNR, the less distortion has been introduced. Ideally, if HDL output and the source image are identical, PSNR outputs **inf**. In our example, this happens only when $gammaValue$ equals 1 (i.e., both Gamma and De-gamma blocks pass the source image through).

We can also use Gamma to corrupt a source image by making it brighter, followed by a De-gamma correction for image recovery.

Generate HDL Code and Verify Its Behavior

To check and generate the HDL code referenced in this example, you must have an HDL Coder™ license.

To generate the HDL code, use the following command:

```
makehdl('GammaCorrectionHDL/Pixel-Stream Gamma Compensation')
```

To infer a RAM to implement a lookup table used in the **Gamma Corrector**, the LUTRegisterResetType property is set to none. To access this property, right click the **Gamma Corrector** block inside **Pixel-Stream Gamma Compensation**, and navigate to HDL Coder -> HDL Block Properties ...

To generate test bench, use the following command:

```
makehdltb('GammaCorrectionHDL/Pixel-Stream Gamma Compensation')
```

Histogram Equalization

This example shows how to use the Vision HDL Toolbox™ Histogram library block to implement histogram equalization.

This example model provides a hardware-compatible algorithm. You can generate HDL code from this algorithm, and implement it on a board using a Xilinx™ Zynq™ reference design. See "Histogram Equalization with Zynq-Based Hardware" (Vision HDL Toolbox Support Package for Xilinx Zynq-Based Hardware).

Introduction

The model shows how to use the Histogram library block to enhance the contrast of images by applying the histogram equalization. To learn more, refer to the Histogram block reference page. There are three components in this histogram equalization example.

- **Video Partition** partitions a big image into four non-overlapping small images for parallel histogram computation.
- **HDLHistogram** computes the accumulated histogram of the image.
- **Equalization** applies the equalized histogram to the original image and generates the contrast-enhanced image.
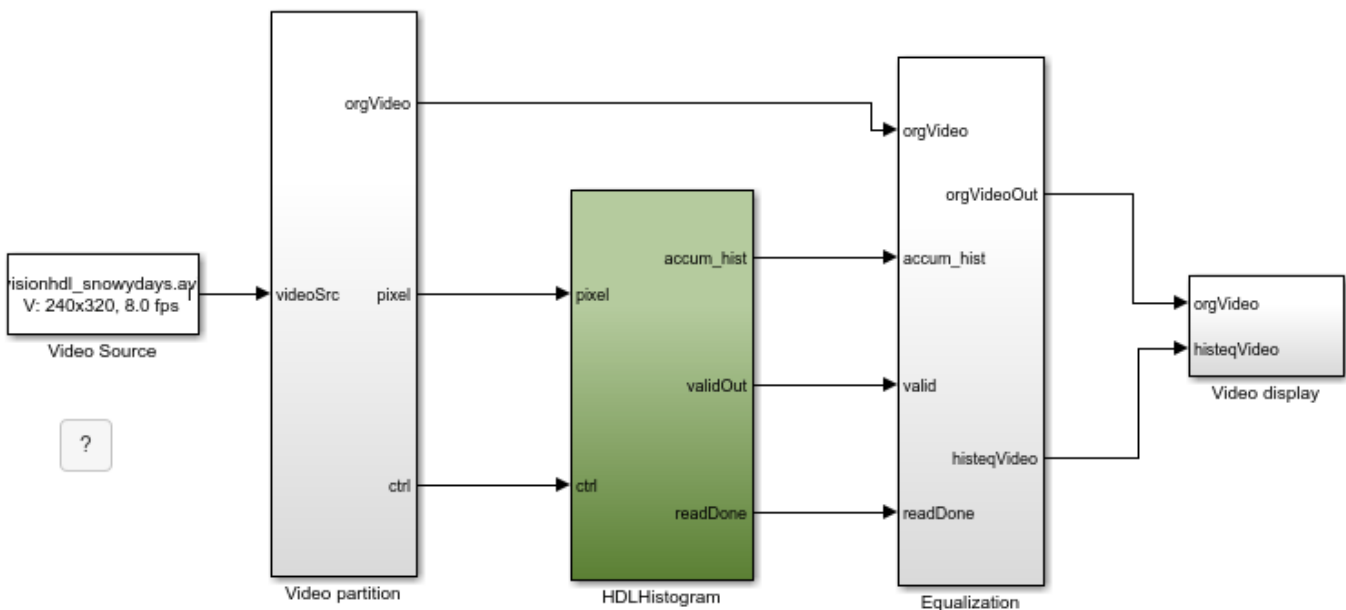
Copyright 2014 The MathWorks, Inc.

Video Partition

There are use cases where histogram is computed over an entire image, or over small regions-of-interest representing sections of the image. Computing histogram of a big image is time consuming. The video partition component in this example divides a big image into four non-overlapping small images. Histogram is computed over the four small images simultaneously. Each input frame is

partitioned into four 120 by 160 small images. Each small image is connected to a Frame To Pixels block to generate pixel streams and corresponding control signals.
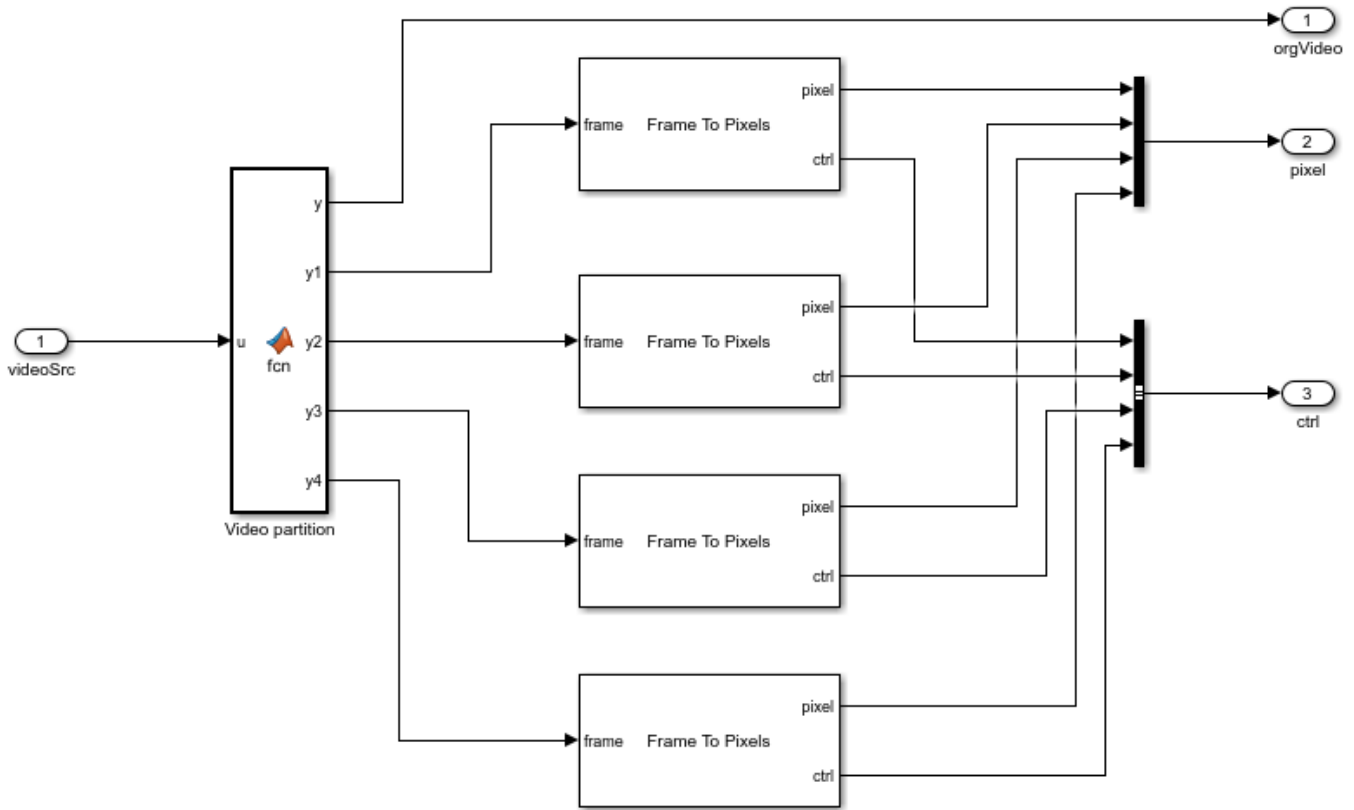
HDLHistogram

HDLHistogram subsystem is optimized for HDL code generation. The histogram of the pixel streams is computed using the Histogram block. Because the input image is grey scale with data type uint8, the input pixels are grouped into 256 bins. The model reads the calculated histogram bins sequentially once the block asserts the *readRdy* signal. The bin values are sent for cumulative histogram calculation. After all 256 bin values are read, the model asserts *binReset* to reset all bins to zero. The collected histogram of each small image is then added together to compute the accumulated histogram of the big image.
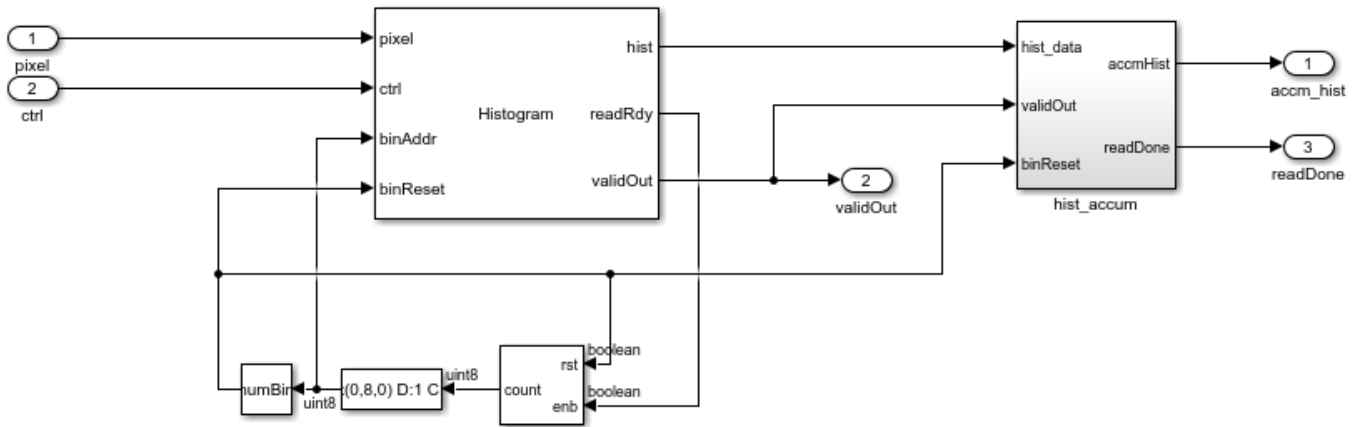
The timing diagram of reading and resetting the histogram bins is shown in the following figure.
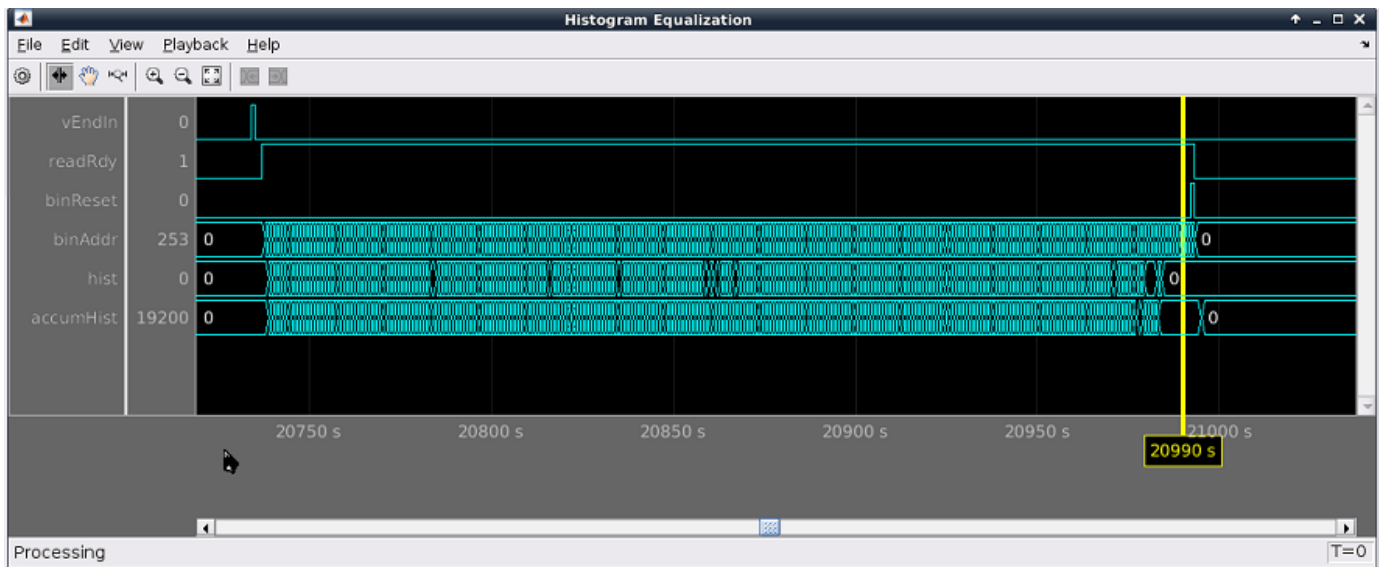
Equalization

Histogram equalization can be applied to the current frame where the accumulated histogram was calculated, or the frame after. If applying to the current frame, the input video needs to be stored. This example delays the input video by one frame and performs uniform equalization to the original video. The equalized video is then compared with the original video.
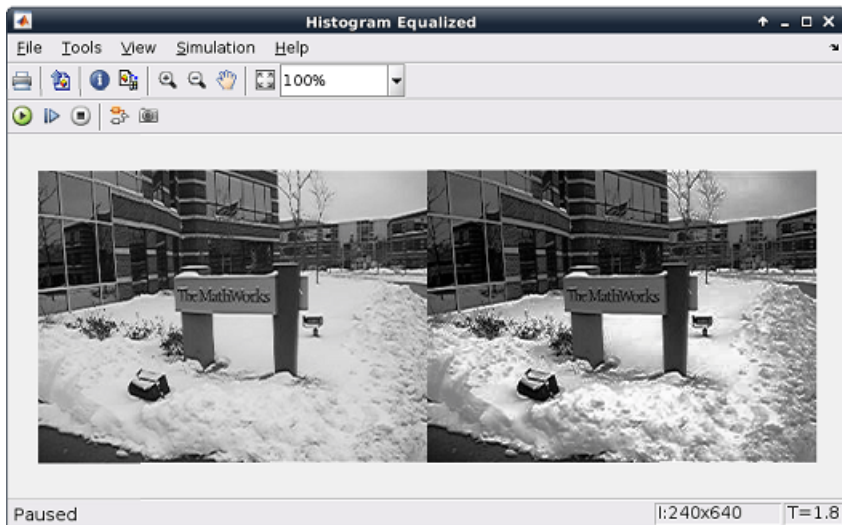
HDL Code Generation

The HDL code generated from the Histogram was synthesized using Xilinx ISE on a Virtex6 (XC6VLX240T-1FFG1156) FPGA, and the circuit ran at about 190 MHz, which is sufficient to process the data in real time.

To check and generate HDL code of this example, you must have an HDL Coder™ license.

You can use the commands

```
makehdl('HistogramEqualizationHDL/HDLHistogram')
```

or

```
makehdltb('HistogramEqualizationHDL/HDLHistogram')
```

to generate HDL code and test bench for the HDLHistogram subsystem. **Note:** Test bench generation takes a long time due to the large data size. Consider reducing the simulation time before generating the test bench.

Edge Detection and Image Overlay

This example shows how to detect and highlight object edges in a video stream. The behavior of the pixel-stream Sobel Edge Detector block, video stream alignment, and overlay, is verified by comparing the results with the same algorithm calculated by the full-frame blocks from the Computer Vision Toolbox™.

This example model provides a hardware-compatible algorithm. You can implement this algorithm on a board using a Xilinx® Zynq® reference design. See "Developing Vision Algorithms for Zynq-Based Hardware" (Vision HDL Toolbox Support Package for Xilinx Zynq-Based Hardware).

Structure of the Example

The EdgeDetectionAndOverlayHDL.slx system is shown below.
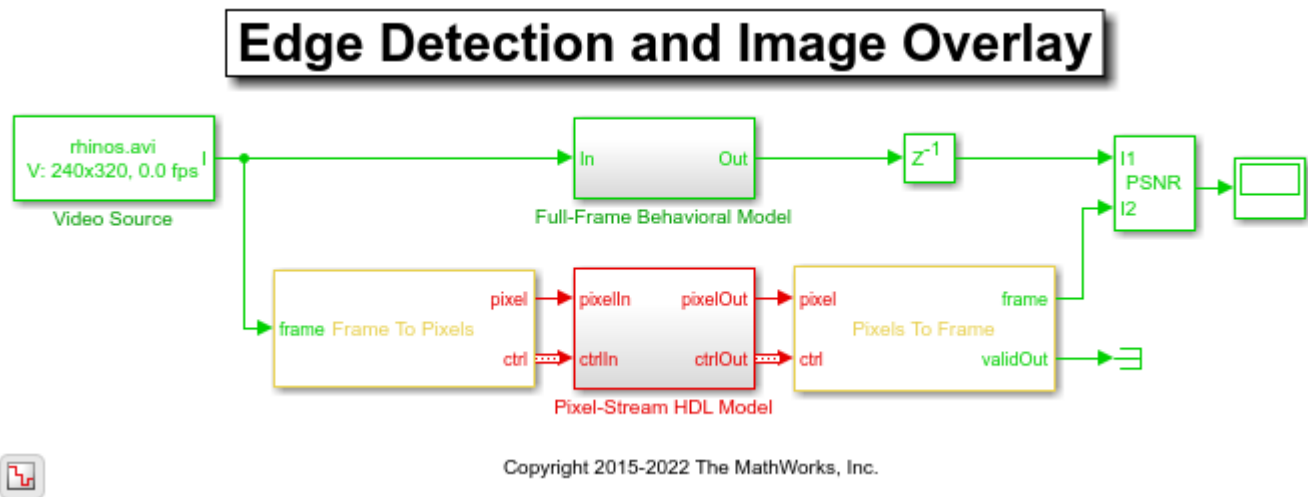
The difference in the color of the lines feeding the **Full-Frame Behavioral Model** and **Pixel-Stream HDL Model** subsystems indicates the change in the image rate on the streaming branch of the model. This rate transition is because the pixel stream is sent out in the same amount of time as the full video frames and therefore it is transmitted at a higher rate.

Full-Frame Behavioral Model

The following diagram shows the structure of the **Full-Frame Behavioral Model** subsystem, which employs the frame-based **Edge Detection** block.
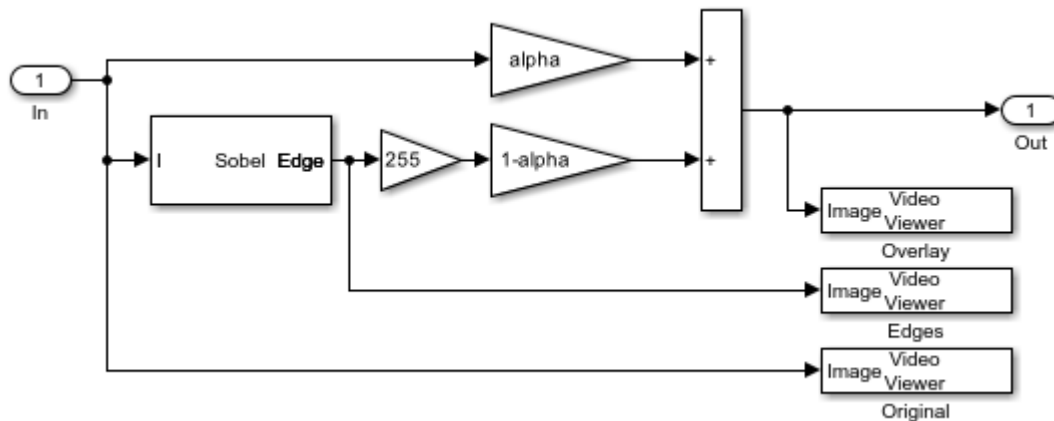
Given that the frame-based **Edge Detection** block does not introduce latency, image overlay is performed by weighting the source image and the **Edge Detection** output image, and adding them together in a straightforward manner.

One frame of the source video, the edge detection result, and the overlaid image are shown from left to right in the diagram below.
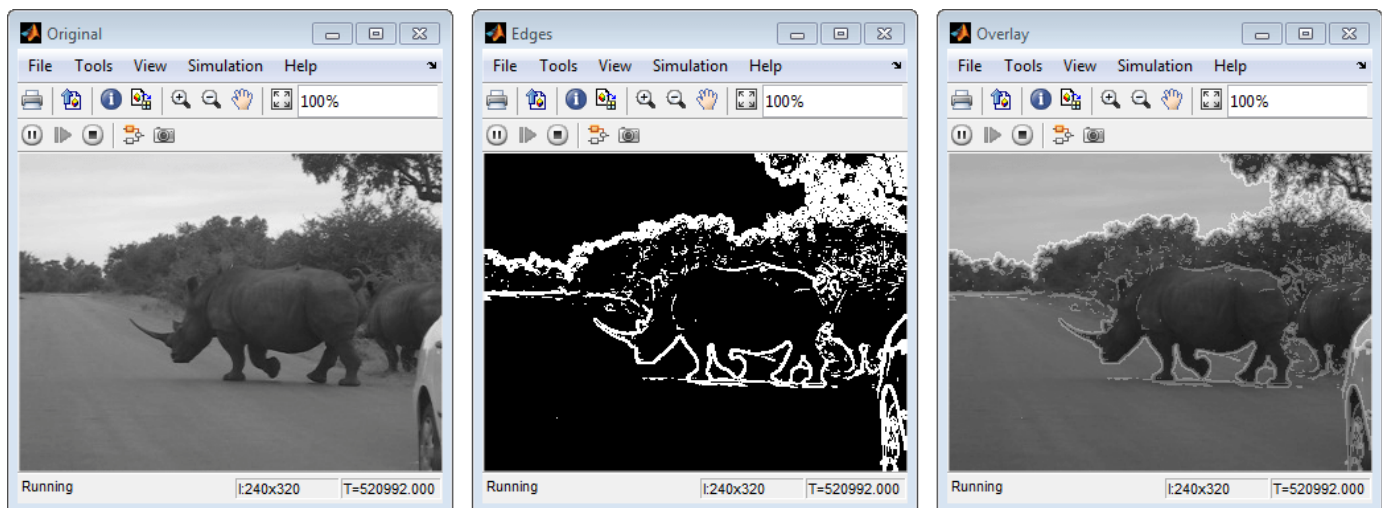
It is a good practice to develop a behavioral system using blocks that process full image frames, the **Full-Frame Behavioral Model** subsystem in this example, before moving forward to working on an FPGA-targeting design. Such a behavioral model helps verify the video processing design. Later on, it can serve as a reference for verifying the implementation of the algorithm targeted to an FPGA. Specifically, the **PSNR** (peak signal-to-noise ratio) block at the top level of the model compares the results from full-frame processing with those from pixel-stream processing.

Frame To Pixels: Generating a Pixel Stream

The task of the **Frame To Pixels** is to convert a full frame image to pixel stream. To simulate the effect of horizontal and vertical blanking periods found in real life hardware video systems, the active image is augmented with non-image data. For more information on the streaming pixel protocol, see "Streaming Pixel Interface" on page 1-2. The **Frame To Pixels** block is configured as shown:

The **Number of components** field is set to 1 for grayscale image input, and the **Video format** field is 240p to match that of the video source.

In this example, the Active Video region corresponds to the 240x320 matrix of the dark image from the upstream **Corruption** block. Six other parameters, namely, **Total pixels per line**, **Total video**

lines, **Starting active line**, **Ending active line**, **Front porch**, and **Back porch** specify how many non-image data will be augmented on the four sides of the Active Video. For more information, see the Frame To Pixels block reference page.

Note that the sample time of the **Video Source** is determined by the product of **Total pixels per line** and **Total video lines**.

Pixel-Stream Edge Detection and Image Overlay

The **Pixel-Stream HDL Model** subsystem is shown in the diagram below. You can generate HDL code from this subsystem.
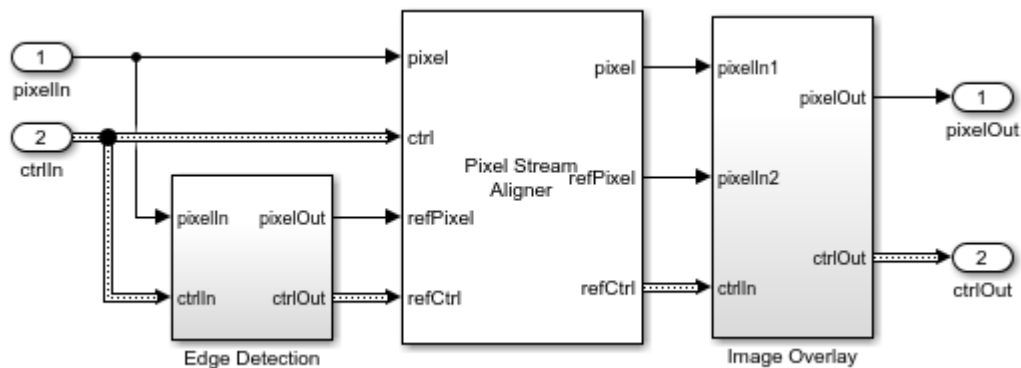
Due to the nature of pixel-stream processing, unlike the **Edge Detection** block in the **Full-Frame Behavioral Model**, the **Edge Detector** block from the Vision HDL Toolbox™ will introduce latency. The latency prevents us from directly weighting and adding two images to obtain the overlaid image. To address this issue, the **Pixel Stream Aligner** block is used to synchronize the two pixel streams before the sum.

To properly use this block, refPixel and refCtrl must be connected to the pixel and control bus that are associated with a delayed pixel stream. In our example, due to the latency introduced by the **Edge Detector**, the pixel stream coming out of the **Edge Detection** subsystem is delayed with respect to that feeding into it. Therefore, the upstream source of refPixel and refCtrl are the pixelOut and ctrlOut signals from the **Edge Detection** subsystem.

Pixels To Frame: Converting Pixel Stream Back to Full Frame

As a companion to **Frame To Pixels** that converts a full image frame to pixel stream, the **Pixels To Frame** block, reversely, converts the pixel stream back to the full frame by making use of the synchronization signals. Since the output of the **Pixels To Frame** block is a 2-D matrix of a full image, there is no need to further carry on the bus containing five synchronization signals.

The **Number of components** field and the **Video format** fields of both Frame To Pixels and Pixels To Frame are set at 1 and 240p, respectively, to match the format of the video source.

Verifying the Pixel Stream Processing Design

While building the streaming portion of the design, the **PSNR** block continuously verifies results against the original full-frame design. The **Delay** block on the top level of the model time-aligns the 2-D matrices for a fair comparison. During the course of the simulation, the **PSNR** block should give **inf** output, indicating that the output image from the **Full-Frame Behavioral Model** matches the image generated from the stream processing **Pixel-Stream HDL Model**.

Exploring the Example

The example allows you to experiment with different threshold and alpha values to examine their effect on the quality of the overlaid images. Specifically, two workspace variables $thresholdValue$ and $alpha$ with initial values 7 and 0.8, respectively, are created upon opening the model. You can modify their values using the MATLAB® command line as follows:

```
thresholdValue=8
alpha=0.5
```

The updated $thresholdValue$ will be propagated to the **Threshold** field of the **Edge Detection** block inside the **Full-Frame Behavioral Model** and the **Edge Detector** block inside **Pixel-Stream HDL Model/Edge Detection**. The $alpha$ value will be propagated to the **Gain1** block in the **Full-Frame Behavioral Model** and **Pixel-Stream HDL Model/Image Overlay**, and the value of $1 - alpha$ goes to **Gain2** blocks. Closing the model clears both variables from your workspace.

In this example, the valid range of $thresholdValue$ is between 0 and 256, inclusive. Setting $thresholdValue$ equal to or greater than 257 triggers a message **Parameter overflow occurred for 'threshold'**. The higher you set the $thresholdValue$, the smaller the amount of edges the example finds in the video.

The valid range of $alpha$ is between 0 and 1, inclusive. It determines the weights for edge detection output image and the original source image before adding them. The overlay operation is a linear interpolation according to the following formula.

```
overlaid image = alpha*source image + (1-alpha)*edge image.
```

Therefore, when $alpha = 0$, the overlaid image is the edge detection output, and when $alpha = 1$ it becomes the source image.

Generate HDL Code and Verify Its Behavior

To check and generate the HDL code referenced in this example, you must have an HDL Coder™ license.

To generate the HDL code, use the following command:

```
makehdl('EdgeDetectionAndOverlayHDL/Pixel-Stream HDL Model');
```

To generate a test bench, use the following command:

```
makehdltb('EdgeDetectionAndOverlayHDL/Pixel-Stream HDL Model');
```

Edge Detection and Image Overlay with Impaired Frame

This example shows how to introduce impairments in order to test a design with imperfect video input.

When designing video processing algorithms, an important concern is the quality of the incoming video stream. Real-life video systems, like surveillance cameras or camcorders, produce imperfect signals. The stream can contain errors such as active lines of unequal length, glitches, or incomplete frames. In simulation, a streaming video source will usually produce perfect signals. When you use the **Frame To Pixels** block from the Vision HDL Toolbox™, all lines are of equal size, and all frames are complete. A video algorithm that simulates well under these conditions does not guarantee its effectiveness on an FPGA that connects to a real-world video source. To assess the robustness of a video algorithm under nonideal real-world video signals, it is practical to introduce impairments in the pixel stream.

This example extends the "Edge Detection and Image Overlay" on page 2-26 example by manually masking off the leading control signals of a frame to resemble a scenario where the algorithm starts in the middle of a frame. Such test scenarios are necessary to prove robustness of streaming video designs.

It is beneficial to go over the "Edge Detection and Image Overlay" on page 2-26 example before proceeding to this example.

Structure of the Example

The structure of this example is shown below, which closely follows the structure of the pixel-stream processing unit of the model in "Edge Detection and Image Overlay" on page 2-26.
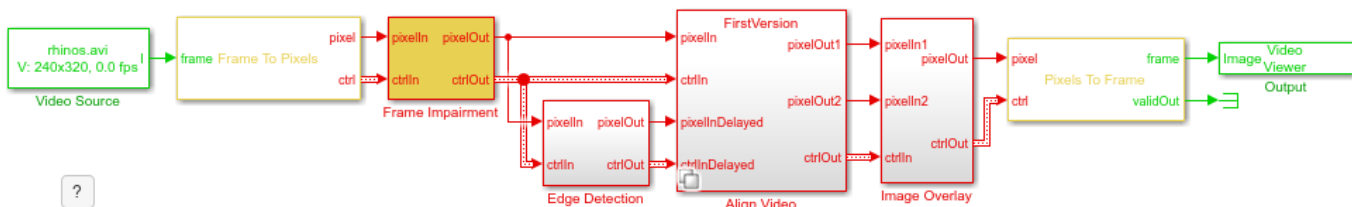
Align Video is implemented as a variant subsystem.
To simulate using the **FirstVersion** variant, type **VERSION=1** at MATLAB prompt.
To simulate using the **SecondVersion** variant, type **VERSION=2** at MATLAB prompt.

Copyright 2014 The MathWorks, Inc.

The **Edge Detection** subsystem implements a Sobel algorithm to highlight the edge of an image. The **Align Video** subsystem is used to synchronize the delayed output of the **EdgeDetector** with the original frame. **Image Overlay** weights and sums up the two time-aligned images.

This material is organized as follows. We first develop an **Align Video** subsystem that works well with perfect video signals. Then, we use the **Frame Impairment** subsystem to mask off the leading control signals of a frame to resemble a scenario where the algorithm starts in the middle of a frame. We will see that such impairment makes **Align Video** ineffective. Finally, a revised version of **Align Video** is developed to address the issue.

Align Video is implemented as a variant subsystem. You can use the variable VERSION in workspace to select which one of the two versions you want to simulate.

Note: Starting in R2017a the Pixel Stream Aligner block replaces the Align Video subsystem shown here. This new block makes setting the line buffer size and number of lines much easier and generates HDL code. In new designs, use the Pixel Stream Aligner block rather than the Align Video subsystem. For an example of how to use the block, see "Edge Detection and Image Overlay" on page 2-26.

First Version of Align Video

The following diagram shows the structure of the first version of the **Align Video** subsystem.
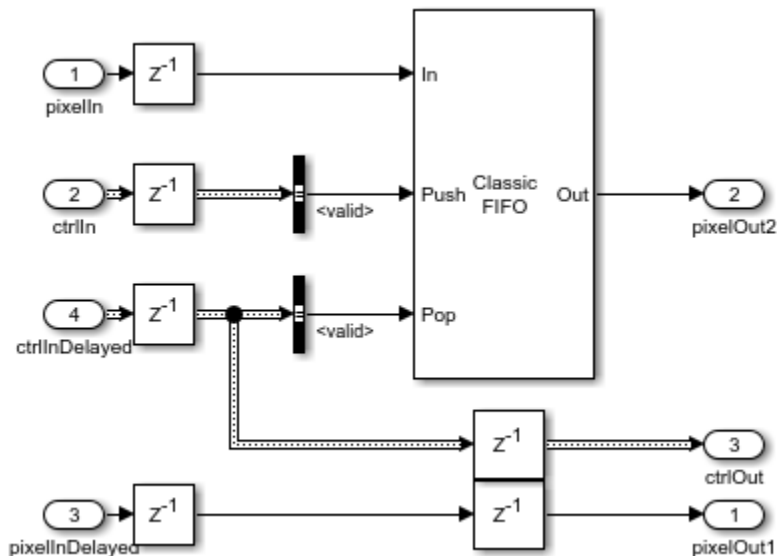
Align Video uses control signals to detect the active region of a frame. For more information on the streaming pixel protocol, see "Streaming Pixel Interface" on page 1-2.

The basic idea of aligning two pixel streams is to buffer valid pixels that come earlier into a FIFO based only on valid signals, and appropriately pop individual pixel from this FIFO based on the valid signal of the delayed pixel-stream.

Test Align Video Using Frame Impairment Subsystem

To illustrate how the **Frame Impairment** subsystem works, consider a 2-by-3 pixel frame. In the figure below, this frame is showed in the dashed rectangle with inactive pixels surrounding it. Inactive pixels include a 1-pixel-wide back porch, a 2-pixel-wide front porch, 1 line before the first active line, and 1 line after the last active line. Both active and inactive pixels are labeled with their grayscale values.
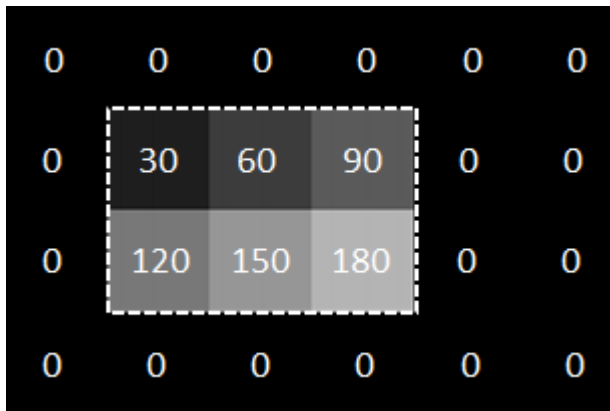
If the **Frame To Pixels** block accepts this 2-by-3 frame as an input and its settings correspond to the porch lengths shown above, then the timing diagram of the **Frame To Pixels** output is illustrated in the upper half of the following diagram.
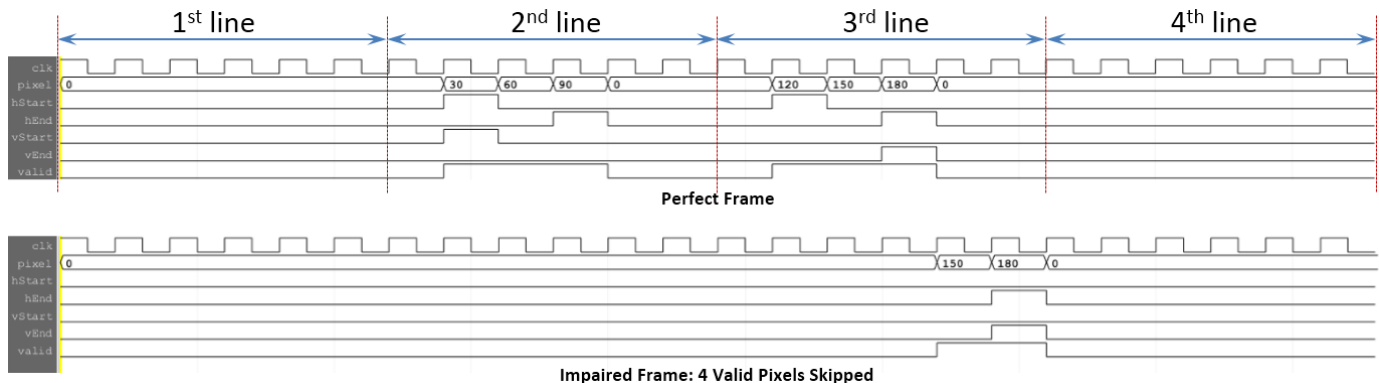
The **Frame Impairment** subsystem skips a configurable number of valid pixels at the beginning of the simulation. For example, if it was configured to skip 4 pixels of the example frame, the result would be as in the lower half of the timing diagram. We can see that by skipping 4 valid pixels, the three valid pixels on the second line (i.e., with intensity values of 30, 60, and 90), and the first valid pixel on the third line, are masked off, along with their associated control signals. Moreover, the **Frame Impairment** subsystem introduces two clock cycle delays. If we enter 0 pixels to skip, it just delays both pixel and ctrl outputs from **Frame To Pixels** by two clock cycles.

Double-click the **Frame Impairment** subsystem and ensure 'Number of valid pixels to skip' is set to 0. As mentioned before, this setting does not impair the frame, all it does is to delay both pixel and ctrl outputs from **Frame To Pixels** by two clock cycles. The output from the video output is shown below, which is expected.
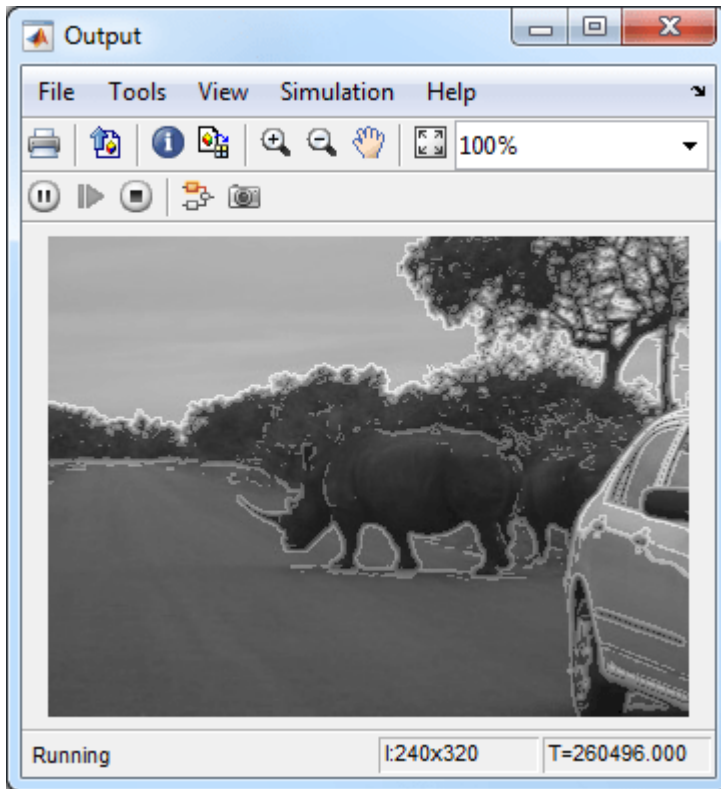
Now, double-click **Frame Impairment** again and enter any positive integer number, say 100, in the 'Number of valid pixels to skip' field.

Rerun the model and the resulting video output is shown below.
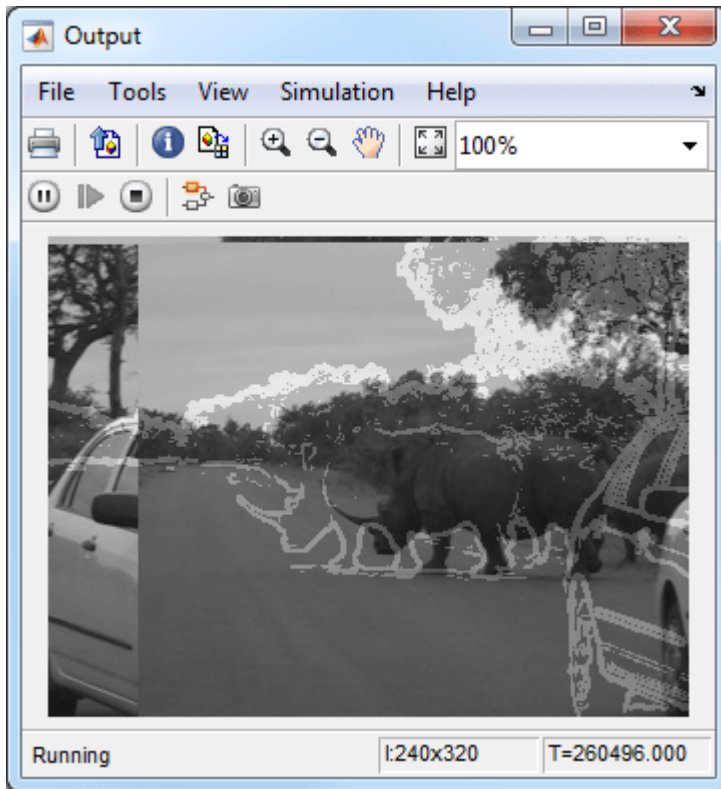
We can see that the edge output is at the right place but the original image is shifted. This output clearly suggests that our first version of **Align Video** is not robust against a pixel stream that starts in the middle of a frame.

Two reasons explain this behavior. Firstly, **EdgeDetector** block starts processing only after seeing a valid frame start, indicated by hStart, vStart, and valid going high at the same clock cycle. The block does not output anything for a partial frame. Secondly, the FIFO, inside the **Align Video** subsystem, starts buffering the frame once the valid signal is true, whether it is a partial frame or a complete frame. Therefore, at the start of the second frame, FIFO has been contaminated with the pixels of the previous partial frame.

Corrected Version of Align Video

Based on the insight gained from the previous section, a revised version of **Align Video** is shown below.
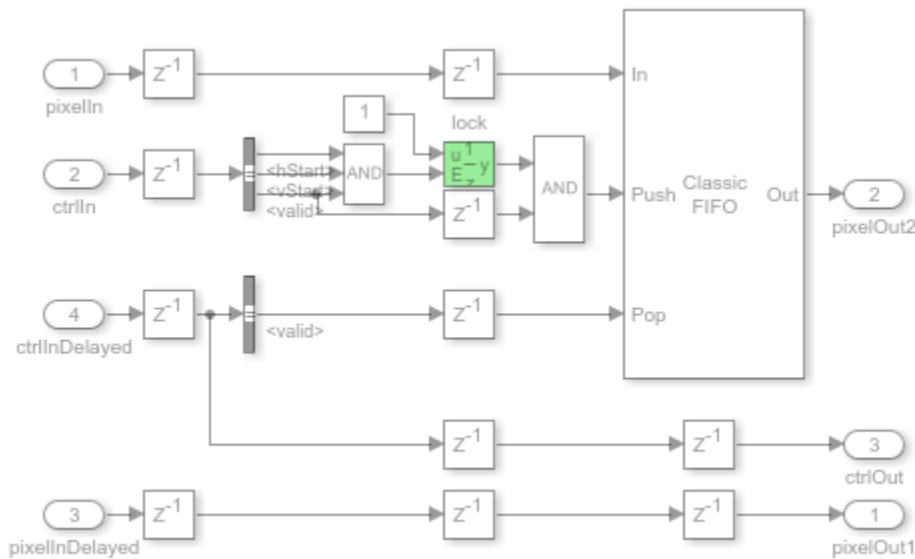
The goal is to only push the pixels of complete frames into the FIFO. If the leading frames are not complete, their valid pixels are ignored.

To achieve this, an enabled register called **lock** is used (highlighted in the diagram above). Its initial value is logical 0. ANDing this 0 with a delayed version of valid always gives logical 0. This prevents any valid pixels from being pushed into FIFO. The **lock** toggles its output from logical 0 to 1 only when hStart, vStart, and valid signals assert high, an indicator of the start of a new frame. After **lock** toggles to 1, the 'push' input of FIFO now follows a delayed version of the valid signal. So the valid pixels of a new frame will be buffered in FIFO.

To test this revised implementation, type the following command at MATLAB prompt.

```
VERSION=2;
```

Rerun the simulation. Now the edge output and the original image are perfectly aligned.

Noise Removal and Image Sharpening

This example shows how to implement a front-end module of an image processing design. This front-end module removes noise and sharpens the image to provide a better initial condition for the subsequent processing.

An object out of focus results in a blurred image. Dead or stuck pixels on the camera or video sensor, or thermal noise from hardware components, contribute to the noise in the image. In this example, the front-end module is implemented using two pixel-stream filter blocks from the Vision HDL Toolbox™. The median filter removes the noise and the image filter sharpens the image. The example compares the pixel-stream results with those generated by the full-frame blocks from the Computer Vision System Toolbox™.

This example model provides a hardware-compatible algorithm. You can implement this algorithm on a board using a Xilinx™ Zynq™ reference design. See "Image Sharpening with Zynq-Based Hardware" (Vision HDL Toolbox Support Package for Xilinx Zynq-Based Hardware).

Structure of the Example

Computer Vision Toolbox blocks operate on an entire frame at a time. Vision HDL Toolbox blocks operate on a stream of pixel data, one pixel at a time. The conversion blocks in Vision HDL Toolbox, Frame To Pixels and Pixels To Frame, enable you to simulate streaming-pixel designs alongside full-frame designs.

The NoiseRemovalAndImageSharpeningHDL.slx system is shown below.
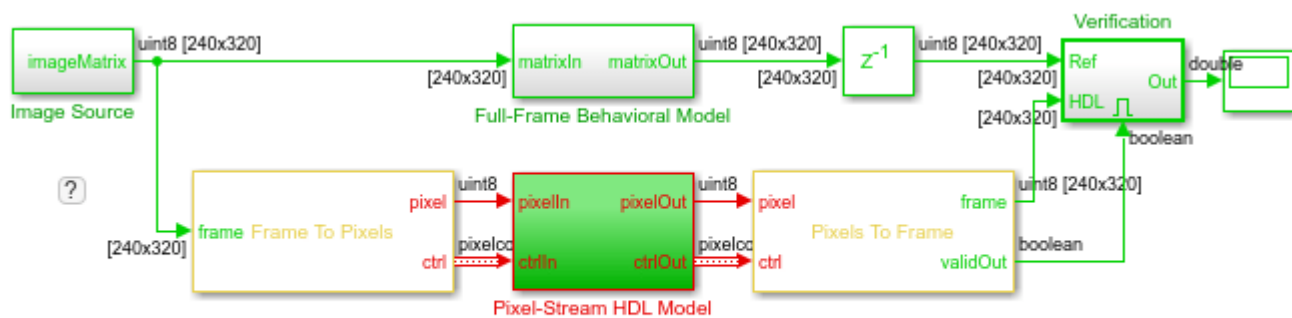
Copyright 2018 The MathWorks, Inc.

The following diagram shows the structure of the Full-Frame Behavioral Model subsystem, which consists of the frame-based Median Filter and 2-D FIR Filter. As mentioned before, median filter removes the noise and 2-D FIR Filter is configured to sharpen the image.
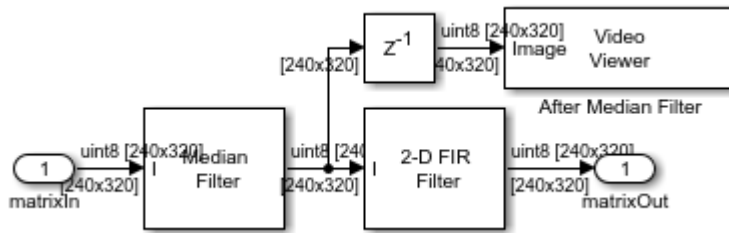
The Pixel-Stream HDL Model subsystem contains the streaming implementation of the median filter and 2-D FIR filter, as shown in the diagram below. You can generate HDL code from the Pixel-Stream HDL Model subsystem.
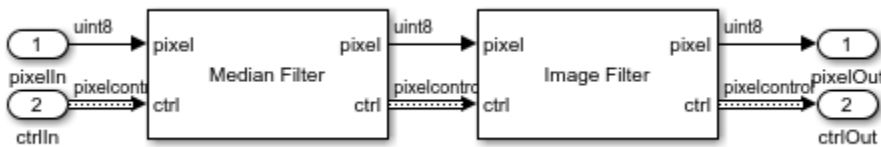
The Verification subsystem compares the results from full-frame processing with those from pixel-stream processing.

One frame of the blurred and noisy source video, its de-noised version after median filtering, and the sharpened output after 2-D FIR filtering, are shown from left to right in the diagram below.
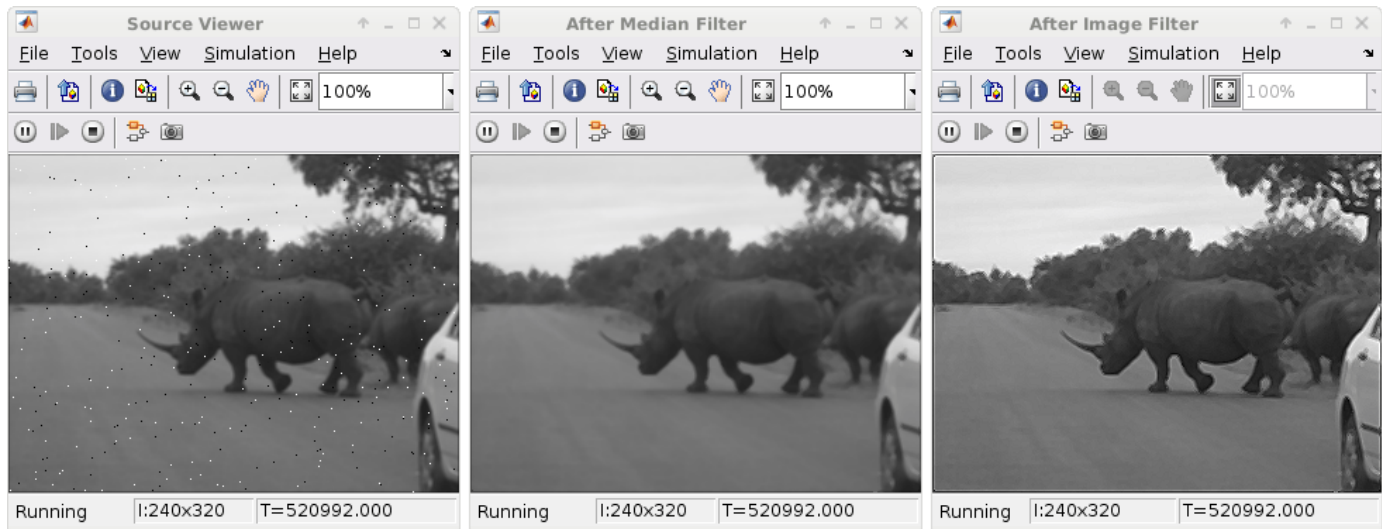
Image Source

The following figure shows the Image Source subsystem.
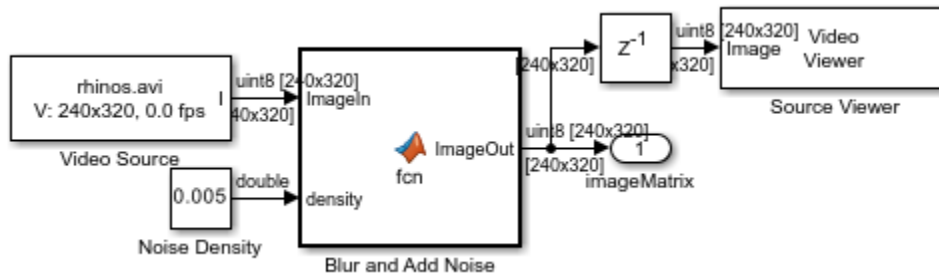
The Image Source block imports a grayscale image, then uses a MATLAB function block named Blur and Add Noise to blur the image and inject salt-and-pepper noise. The IMFILTER function uses a 3-by-3 averaging kernel to blur the image. The salt-and-pepper noise is injected by calling the IMNOISE(I,'salt & pepper',D) command, where D is the noise density defined as the ratio of the combined number of salt and pepper pixels to the total pixels in the image. This density value is specified by the Noise Density constant block, and it must be between 0 and 1. The Image Source subsystem outputs a 2-D matrix of a full image.

Frame To Pixels: Generating a Pixel Stream

The Frame To Pixels block converts a full image frame to a pixel stream. The Number of components field is set to 1 for grayscale image input, and the Video format field is 240p to match that of the video source. The sample time of the Video Source is determined by the product of Total pixels per line and Total video lines in the Frame To Pixels block. For more information, see the Frame To Pixels block reference page.

Pixel-Stream HDL Model

The Median Filter block is used to remove the salt and pepper noise. To learn more, refer to the Median Filter block reference page.

Based on the filter coefficients, the Image Filter block can be used to blur, sharpen, or detect the edges of the recovered image after median filtering. In this example, Image Filter is configured to sharpen an image. To learn more, refer to the Image Filter block reference page.

Pixels To Frame: Converting Pixel Stream Back to Full Frame

The Pixels To Frame block converts a pixel stream to the full frame by making use of the synchronization signals. The Number of components field and the Video format field of the Pixels To Frame are set at 1 and 240p, respectively, to match the format of the video source.

Verifying the Pixel-Stream Processing Design

The Verification subsystem, as shown below, verifies the results from the pixel-stream HDL model against the full-frame behavioral model.
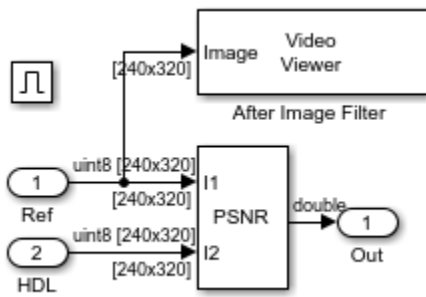
The peak signal to noise ratio (PSNR) is calculated between the reference image and the stream processed image. Ideally, the ratio should be inf, indicating that the output image from the Full-Frame Behavioral Model matches that generated from the Pixel-Stream HDL Model.

Generate HDL Code and Verify Its Behavior

To check and generate the HDL code referenced in this example, you must have an HDL Coder™ license.

To generate the HDL code, use the following command:

```
makehdl('NoiseRemovalAndImageSharpeningHDL/Pixel-Stream HDL Model');
```

To generate test bench, use the following command:

```
makehdltb('NoiseRemovalAndImageSharpeningHDL/Pixel-Stream HDL Model');
```

Multi-Zone Metering

This example shows how to use the Image Statistics block to perform multi-zone metering to extract a region of interest (ROI).

There are numerous applications where the input video is divided into several zones, and the statistic is then computed over each zone. For example, many auto-exposure algorithms compute the difference in the mean intensity between zones. This allows the shutter controller logic to determine whether the image is under-exposed (overall low illumination), correctly-exposed (uniform illumination) or over-exposed (one or more ROIs have a larger mean).

Introduction

The MultizoneMeteringHDL.slx system is shown below.
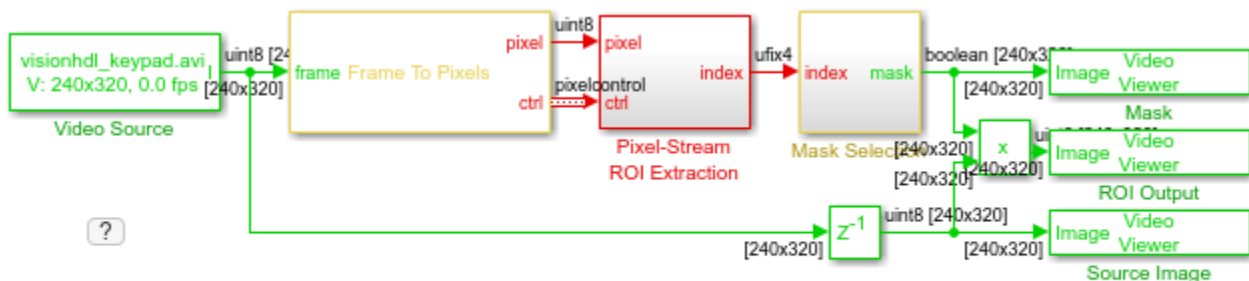
Copyright 2015 The MathWorks, Inc.

The green and red lines represent full-frame processing and pixel-stream processing, respectively. The color difference indicates the change in the image rate on the streaming branch of the model. This rate transition is because the pixel stream is sent out in the same amount of time as the full video frames and therefore it is transmitted at a higher rate.

In this example, the **Pixel-Stream ROI extraction** subsystem calculates the mean intensity value over 12 predefined ROIs in a frame and outputs the index number (1-12) that corresponds to the most illuminated ROI. The downstream **Mask Selection** subsystem accepts this index number and outputs the associated binary mask image. The binary mask image is applied to the source video to display only the most illuminated ROI, and mask off the other 11 ROIs. The **Delay** block at the top level of the model is used to match the latency introduced by pixel-stream processing.

One frame of the source image, the binary mask image, and the ROI output, are shown from left to right in the diagram below.
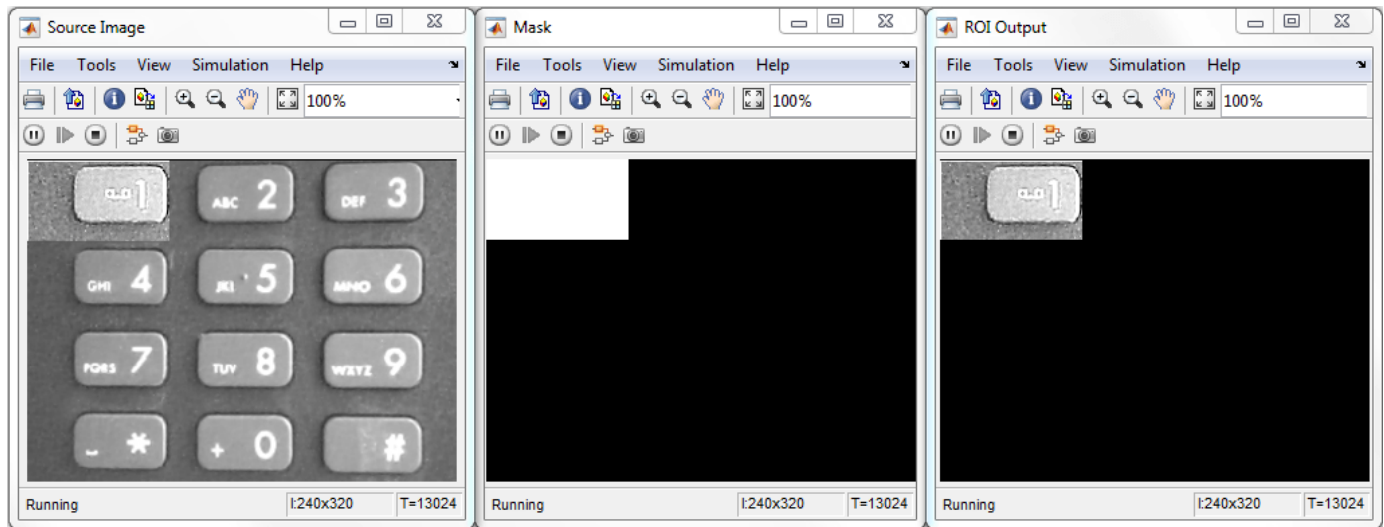
You can generate HDL code from the **Pixel-Stream ROI Extraction** subsystem.

Video Source

The video format is 240p. Each frame consists of 240 lines and 320 pixels per line. In this example, video frames are divided into 12 non-overlapping rectangular ROIs, denoted as ROI number 1 to 12, as shown in the diagram below. Each ROI includes one key of the input keypad image.

ROI number 1 has a 107-pixel width and a 60-pixel height, and the (x,y) coordinate of its top-left pixel is (1,1). ROI number 2 has a 107-pixel width and a 60-pixel height, and the coordinate of its top left pixel is (108,1), and so on. The first frame of the input video has brighter pixels within ROI number 1, as shown above. The second frame has brighter pixels within ROI number 2, and so on.

Frame To Pixels: Generating a Pixel Stream

Frame To Pixels converts a full-frame image to a pixel stream. To simulate the effect of horizontal and vertical blanking periods found in real life hardware video systems, the active image is

augmented with non-image data. For more information on the streaming pixel protocol, see "Streaming Pixel Interface" on page 1-2. The **Frame To Pixels** block is configured as shown:

The **Number of components** field is set to 1 for grayscale image input, and the **Video format** field is 240p to match that of the video source.

In this example, the Active Video region corresponds to the 240x320 matrix of the source image. Six other parameters, namely, **Total pixels per line**, **Total video lines**, **Starting active line**, **Ending active line**, **Front porch**, and **Back porch** specify how many non-image data will be augmented on the four sides of the Active Video. For more information, see the Frame To Pixels block reference page.

Note that the sample time of the **Video Source** is determined by the product of **Total pixels per line** and **Total video lines**.

Pixel-Stream ROI Extraction

The **Pixel-Stream ROI Extraction** subsystem contains two subsystems, namely, **Multi-Zone Metering** and **ROI Indexer**.
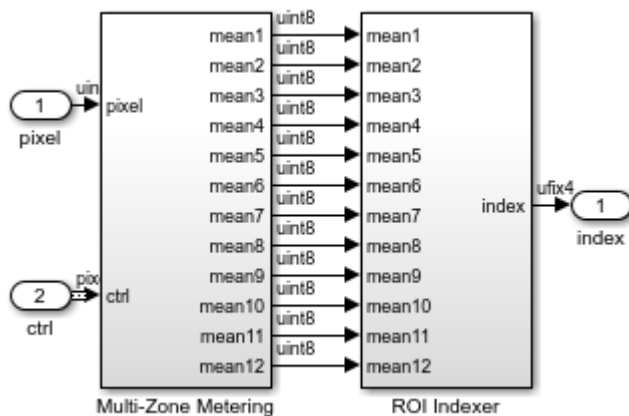
The **Multi-Zone Metering** subsystem computes the mean intensity value over the 12 predefined ROIs. The resulting 12 mean values are passed to the downstream **ROI Indexer** subsystem. **ROI Indexer** outputs the index (1-12) of the ROI that has the maximum mean intensity value (or equivalently, the most illuminated ROI) among the 12 candidates.

The structure of the **Multi-Zone Metering** subsystem is shown in the diagram below.
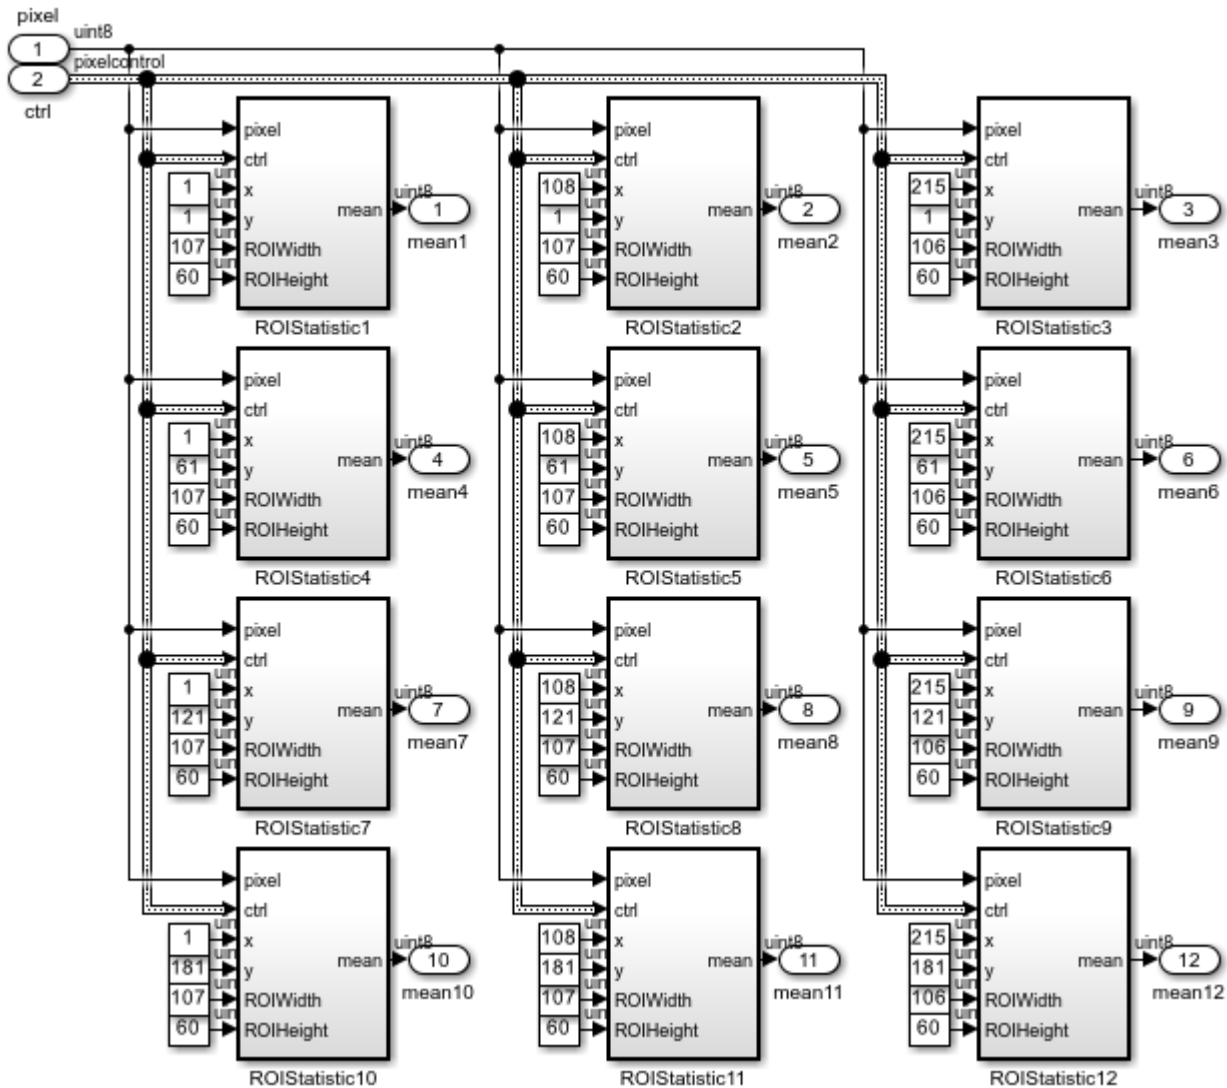
The **Multi-Zone Metering** subsystem contains 12 identical **ROIStatistic** subsystems. Each instance of **ROIStatistic** calculates the mean intensity value over one ROI. All of the 12 **ROIStatistic** subsystems take pixel and ctrl as their first two inputs. The remaining four inputs specify which ROI this subsystem works on and they are different from one subsystem to another. For example, the **ROIStaticstic1** subsystem focuses on ROI number 1 by accepting the (x,y) coordinate of the top left pixel (1,1), ROI width of 107, and height 60. Similarly, the **ROIStaticstic12** subsystem focuses on ROI number 12, whose (x,y) coordinate of the top left pixel is (215,181), and whose width and height are 106 and 60, respectively.

The **ROIStatistic1** - **ROIStatistic12** subsystems share the same structure shown below.
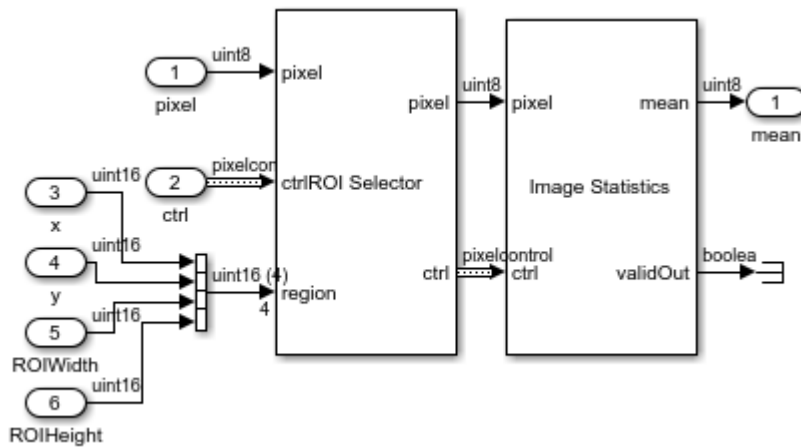
It contains a **ROI Selector** block followed by an **Image Statistics** block. The **ROI Selector** block manipulates the control signal of the original 240p image, and constructs the control signals associated only with the ROI specified by (x,y) pair, ROIWidth, and ROIHeight.

Mask Selection

The structure of the **Mask Selection** subsystem is shown below.
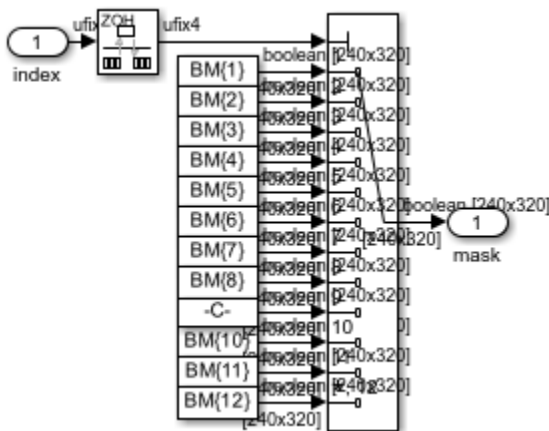
Twelve mask images are available, corresponding to the 12 different ROIs. These mask patterns are shown as BM{1} to BM{12} in the above diagram. When you open the model, the model loads the predefined BM cell array into the workspace. Masks are binary images with 240p video format. For mask BM{n} (n=1,2,...,12), the ROI number n is filled with logical 1 pixels (white) and all the other 11 ROIs are filled with logical 0 pixels (black). Based on the index input (1-12), the **Mask Selection** subsystem outputs the associated binary mask image.

HDL Code Generation

To check and generate the HDL code referenced in this example, you must have an HDL Coder™ license.

To generate the HDL code, use the following command.

```
makehdl('MultizoneMeteringHDL/Pixel-Stream ROI Extraction')
```

To generate a test bench, use the following command. Note that the test bench generation takes a long time due to the large data size. You may want to reduce the simulation time before generating the test bench.

```
makehdltb('MultizoneMeteringHDL/Pixel-Stream ROI Extraction')
```

Harris Corner Detection

This example shows how to use edge detection as the first step in corner detection. The algorithm is suitable for FPGAs.

Corner detection is used in computer vision systems to find features in an image. It is often one of the first steps in applications like motion detection, tracking, image registration and object recognition.

A corner is intuitively defined as the intersection of two edges. This example uses the Harris & Stephens algorithm [1] in which the computation is simplified using an approximation of the eigenvalues of the Harris matrix. For another corner detection algorithm for FPGAs, see the "FAST Corner Detection" on page 2-53 example.

This example model provides a hardware-compatible algorithm. You can implement this algorithm on a board using a Xilinx™ Zynq™ reference design. See "Corner Detection and Image Overlay with Zynq-Based Hardware" (Vision HDL Toolbox Support Package for Xilinx Zynq-Based Hardware).

Introduction

The CornerDetectionHDL.slx system is shown below. The HDL Corner Algorithm subsystem contains a Corner Detector block with the **Method** parameter set to `Harris`.
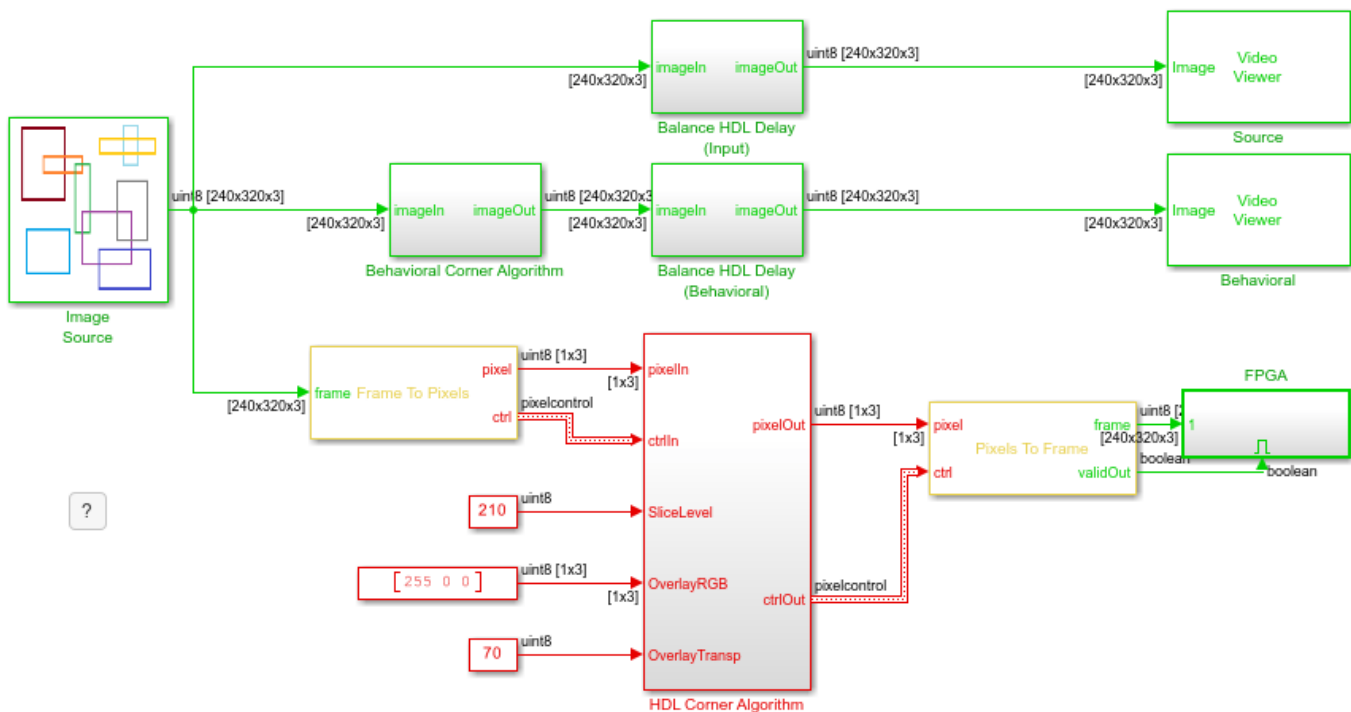
Copyright 2017 The MathWorks, Inc.

First Step: Find the Gradients

The first step in the Harris algorithm is to find the edges in the image. The Corner Detector block uses two gradient image filters with coefficients $\begin{bmatrix} 1 & 0 & -1 \end{bmatrix}$ and $\begin{bmatrix} 1 \\ 0 \\ -1 \end{bmatrix}$ to produce gradients G_x and G_y. Square and cross-multiply to form G_x^2, G_y^2 and G_{xy}.

Second Step: Circular Filtering

The second step of the algorithm is to perform Gaussian filtering to average G_x^2, G_y^2 and G_{xy} over a circular window. The size of the circular window determines the scale of the detected corner. The block uses a 5x5 window. For three components, the block uses three filters with the same filter coefficients.

Final Step: Form the Harris Matrix

The final step of the algorithm is to estimate the eigenvalue of the Harris matrix. The Harris matrix is a symmetric matrix similar to a covariance matrix. The main diagonal is composed of the two averages of the gradients squared $\langle G_x^2 \rangle$ and $\langle G_y^2 \rangle$. The off diagonal elements are the averages of the gradient cross-product $\langle G_{xy} \rangle$. The Harris matrix is:

$$A_{Harris} = \begin{bmatrix} \langle G_x^2 \rangle & \langle G_{xy} \rangle \\ \langle G_{xy} \rangle & \langle G_y^2 \rangle \end{bmatrix}$$

Compute the Response from the Harris Matrix

The key simplification of the Harris algorithm is estimating the eigenvalues of the Harris matrix as the determinant minus the scaled trace squared.

$R = det(A_{Harris}) - k \cdot Tr^2(A_{Harris})$ where k is a constant typically 0.04.

The corner metric response, R, expressed using the gradients is:

$$R = \left(\langle G_x^2 \rangle \cdot \langle G_y^2 \rangle - \langle G_{xy} \rangle^2 \right) - k \cdot \left(\langle G_x^2 \rangle + \langle G_y^2 \rangle \right)^2$$

When the response is larger than a predefined threshold, a corner is detected:

$$R \quad > \quad k_{thresh}$$

$$\left(\langle G_x^2 \rangle \cdot \langle G_y^2 \rangle - \langle G_{xy} \rangle^2 \right) - k \cdot \left(\langle G_x^2 \rangle + \langle G_y^2 \rangle \right)^2 \quad > \quad k_{thresh}$$

Fixed-Point Settings

The overall function from input image to output corner metric response is a fourth-order polynomial. This leads to some challenges determining the fixed-point scaling for each step of the computation. Since we are targeting FPGAs with built-in multipliers, the best strategy is to allow bit growth until the multiplier size is reached and then start to quantize results on a selective basis to stay within the bounds of the provided multipliers.

The input pixel stream is 8-bit grayscale pixel data. Computing the gradients does not add much bit-growth since the filter kernel has only +1 and -1 coefficients. The result is a full-precision 9-bit signed fixed-point type.

Squaring and cross-multiplying the gradients produces signed 18-bit results, still in full precision. Many common FPGA multipliers have 18-bit or 20-bit input wordlengths, so you will have to quantize at the next step.

The next step is to apply a circular window to the three components using three Image Filters with Gaussian coefficients. The coefficients are quantized to 18-bit unsigned numbers to fit the FPGA multipliers. To find the best fraction precision for the coefficients, create a fixed-point number using the fi() function but only specifying the wordlength. In this case a fractional scaling of 21-bits is best since the largest value in the coefficient matrix is between 1/8 and 1/16.

```
coeffs = fi(fspecial('gaussian',[5,5],1.5),0,18)


coeffs =

    0.0144    0.0281    0.0351    0.0281    0.0144
    0.0281    0.0547    0.0683    0.0547    0.0281
    0.0351    0.0683    0.0853    0.0683    0.0351
    0.0281    0.0547    0.0683    0.0547    0.0281
    0.0144    0.0281    0.0351    0.0281    0.0144

          DataTypeMode: Fixed-point: binary point scaling
           Signedness: Unsigned
           WordLength: 18
        FractionLength: 21
```
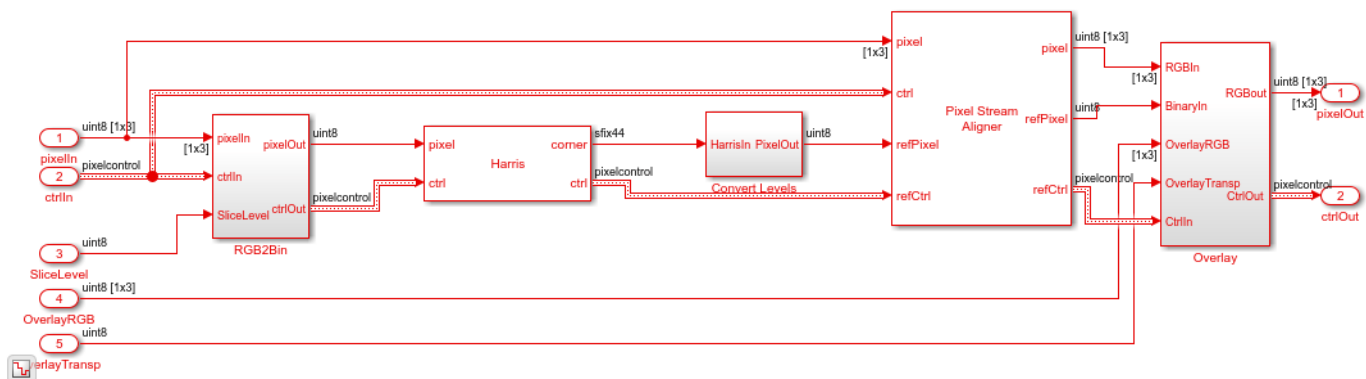

Results of the Simulation

You can see that the resulting images from the simulation are very similar but not exactly the same. The small differences in simulation results are because the behavioral model uses C integer arithmetic rules and the quantization is different from the HDL-ready corner detection block.

Using Simulink, you can understand these differences and decide if the errors are allowable for your application. If they are not acceptable, you can increase the bit-widths of the operators, although this increases the area used in the FPGA.

HDL Code Generation

To check and generate the HDL code referenced in this example, you must have an HDL Coder™ license.

To generate the HDL code, use the following command.

```
makehdl('CornerDetectionHDL/HDL Corner Algorithm')
```

To generate the test bench, use the following command. Note that test bench generation takes a long time due to the large data size. You may want to reduce the simulation time before generating the test bench.

```
makehdltb('CornerDetectionHDL/HDL Corner Algorithm')
```

The part of this model that you can implement on an FPGA is the part between the Frame To Pixels and Pixels To Frame blocks. That is the subsystem called HDL Corner Algorithm, which includes all elements of the corner detection algorithm seen above. The rest of the model, including the Behavioral Corner Algorithm and the sources and sinks, form our Simulink test bench.

Going Further

The Harris & Stephens algorithm is based on approximating the eigenvalues of the Harris matrix as shown above. The Harris algorithm uses $R = det(A_{Harris}) - k \cdot Tr^2(A_{Harris})$ as a metric, avoiding any division or square-root operations. Another way to do corner detection is to compute the actual eigenvalues.

The analytical solution for the eigenvalues of a 2x2 matrix is well-known and can also be used in corner detection. When the eigenvalues are both positive and large with the same scale, a corner has been found.

$$\lambda_1 = \frac{Tr(A)}{2} + \sqrt{\frac{Tr^2(A)}{4} - det(A)}$$

$$\lambda_2 = \frac{Tr(A)}{2} - \sqrt{\frac{Tr^2(A)}{4} - det(A)}$$

Substituting in our A_{Harris} values we get:

$$\lambda_1 = \left(\frac{\langle G_x^2 \rangle + \langle G_y^2 \rangle}{2}\right) + \sqrt{\left(\frac{\langle G_x^2 \rangle + \langle G_y^2 \rangle}{2}\right)^2 - \left(\langle G_x^2 \rangle \cdot \langle G_y^2 \rangle - \langle G_{xy} \rangle^2\right)}$$

$$\lambda_2 = \left(\frac{\langle G_x^2 \rangle + \langle G_y^2 \rangle}{2}\right) - \sqrt{\left(\frac{\langle G_x^2 \rangle + \langle G_y^2 \rangle}{2}\right)^2 - \left(\langle G_x^2 \rangle \cdot \langle G_y^2 \rangle - \langle G_{xy} \rangle^2\right)}$$

For FPGA implementation it is important to notice the repeated value of $\frac{Tr(A)}{2}$. We can compute this value once and then square to combine with $det(A)$. This means that the eigenvalue algorithm requires only two multipliers but at the expense of more adders and subtractors and a square-root function, which requires several multipliers on its own.

You must then compare both eigenvalues to a constant value to make sure they are large. Since the eigenvalues scale up with image intensity, you also need to make sure they are both around the same size. You can do this by subtracting one from another and making sure that result is smaller than some predefined threshold value. Notice that in this subtraction, the first terms cancel out and you are left with:

$$\lambda_1, \lambda_2 \quad > \quad k_{minimum}$$

$$\lambda_1 - \lambda_2 \quad < \quad k_{thresh}$$

$$2\sqrt{\frac{Tr^2(A)}{4} - det(A)} \quad < \quad k_{thresh}$$

$$\frac{Tr^2(A)}{4} - det(A) \quad < \quad \left(\frac{k_{thresh}}{2}\right)^2$$

You can rearrange this so that it is very similar to Harris metric R above:

$$det(A) - \frac{Tr^2(A)}{4} \quad \geq \quad \left(\frac{k_{thresh}}{2}\right)^2$$

Expanding the matrix gives:

$$\left(\langle G_x^2\rangle \cdot \langle G_y^2\rangle - \langle G_{xy}\rangle^2\right) - \left(\frac{\langle G_x^2\rangle + \langle G_y^2\rangle}{2}\right)^2 \quad \geq \quad \left(\frac{k_{thresh}}{2}\right)^2$$

The similarity between the difference of the eigenvalues and the Harris R metric shows how the Harris approximation works. If you rearrange the terms under the square-root and swap the signs so the result must be greater than or equal to a predefined threshold, you arrive at essentially the Harris metric with some scaling.

References

[1] C. Harris and M. Stephens (1988). "A combined corner and edge detector". Proceedings of the 4th Alvey Vision Conference. pp. 147-151.

FAST Corner Detection

This example shows how to perform corner detection using the features-from-accelerated-segment test (FAST) algorithm. The algorithm is suitable for FPGAs.

Corner detection is used in computer vision systems to find features in an image. It is often one of the first steps in applications like motion detection, tracking, image registration and object recognition.

The FAST algorithm determines if a corner is present by testing a circular area around the potential center of the corner. The test detects a corner if a contiguous section of pixels are either brighter than the center plus a threshold or darker than the center minus a threshold. For another corner detection algorithm for FPGAs, see the "Harris Corner Detection" on page 2-48 example.

In a software implementation the FAST algorithm allows for a quick test to rule out potential corners by only testing the four pixels along the axes. Software algorithms only perform the full test if the quick test passes. A hardware implementation can easily perform all the tests in parallel so a quick test is not particularly advantageous and is not included in this example.

The FAST algorithm can be used at many sizes or scales. This example detects corners using a sixteen-pixel circle. In these sixteen pixels, if any nine contiguous pixel meet the brighter or darker limit then a corner is detected.

MATLAB FAST Corner Detection

The Computer Vision System Toolbox™ includes a software FAST corner detection algorithm in the `detectFASTFeatures` function. This example uses this function as the behavioral model to compare against the FAST algorithm design for hardware in Simulink®. The function has parameters for setting the minimum contrast and the minimum quality.

The minimum contrast parameter is the threshold value that is added or subtracted from the center pixel value before comparing to the ring of pixels.

The minimum quality parameter controls which detected corners are "strong" enough to be marked as actual corners. The strength metric in the original FAST paper is based on summing the differences of the pixels in the circular area to the central pixel [2]. Later versions of this algorithm use a different strength metric based on the smallest change in pixel value that would make the detection no longer a corner. `detectFastFeatures` uses the smallest-change metric.

This code reads the first frame of video, converts it to gray scale, and calls `detectFASTFeatures`. The result is a vector of corner locations. To display the corner locations, use the vector to draw bright green dots over the corner pixels in the output frame.

```
v = VideoReader('rhinos.avi');
I = rgb2gray(readFrame(v));
% create output RGB frame
Y = repmat(I,[1 1 3]);
corners = detectFASTFeatures(I,'minContrast',15/255,'minQuality',1/255);
locs = corners.Location;
for ii = 1:size(locs,1)
    Y(floor(locs(ii,2)),floor(locs(ii,1)),2) = 255; % green dot
end
imshow(Y)
```

Limitations of the FAST Algorithm

Other corner detection methods work very differently from the FAST method and a surprising result is that FAST does not detect corners on computer generated images that are perfectly aligned to the x and y axes. Since the detected corner must have a ring of darker or lighter pixel values around the center that includes both edges of the corner, crisp images do not work well. For example, try the FAST algorithm on the input image used in the Harris "Harris Corner Detection" on page 2-48 example.

```
I = imread('cornerboxes.png');
Ig = rgb2gray(I);
corners = detectFASTFeatures(Ig,'minContrast',15/255,'minQuality',1/255)
```

```
corners =

  0x1 cornerPoints array with properties:

    Location: [0x2 single]
      Metric: [0x1 single]
       Count: 0
```

You can see that the function detected zero corners. This because the FAST algorithm requires a ring of contrasting pixels more than halfway around the center of corner. In the computer generated image, both edges of a box at a corner are in the ring of pixel used, so the test for a corner fails. A work-around to this problem is to add blur (by applying a Gaussian filter) to the image so that the corners are less precise but can be detected. After blurring, the FAST algorithm now detects over 100 corners.

```
h = fspecial('gauss',5);
Ig = imfilter(Ig,h);
corners = detectFASTFeatures(Ig,'minContrast',15/255,'minQuality',1/255)
```

```
locs = corners.Location;
for ii = 1:size(locs,1)
    I(floor(locs(ii,2)),floor(locs(ii,1)),2) = 255; % green dot
end
imshow(I)
```

```
corners =

  136x1 cornerPoints array with properties:

    Location: [136x2 single]
      Metric: [136x1 single]
       Count: 136
```
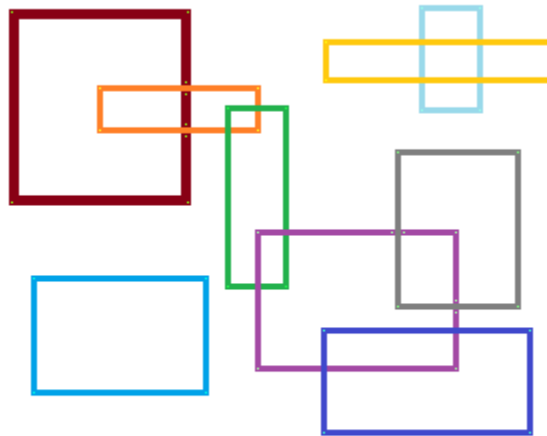

Behavioral Model for Verification

The Simulink model uses the `detectFASTFeatures` function as a behavioral model to verify the results of the hardware algorithm. You can use a MATLAB Function block to run MATLAB code in Simulink.

```
modelname = 'FASTCornerHDL';
open_system(modelname);
set_param(modelname,'SampleTimeColors','on');
set_param(modelname,'SimulationCommand','Update');
set_param(modelname,'Open','on');
set(allchild(0),'Visible','off');
```
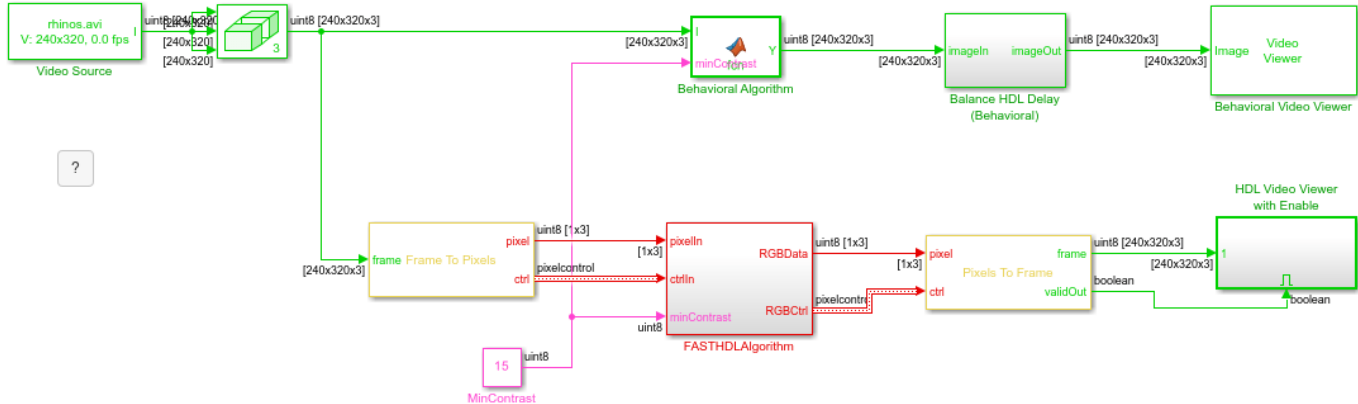
Copyright 2018 The MathWorks, Inc.

The code in a MATLAB Function block must either generate C code or be declared extrinsic. An extrinsic declaration allows the specified function to run in MATLAB while the rest of the MATLAB Function block runs in Simulink. The `detectFASTFeatures` function does not support code generation, so the MATLAB Function block must use an extrinsic helper function.

For frame-by-frame visual comparison, and the ability to vary the contrast parameter, the helper function takes an input image and the minimum contrast as inputs. It returns an output image with green dots marking the detected corners.

```
function Y = FASTHelper(I,minContrast)
Y = I;
corners = detectFASTFeatures(I(:,:,1),'minContrast',double(minContrast)/255,'minQuality',1/255);
locs = corners.Location;
for ii = 1:size(locs,1)
    Y(floor(locs(ii,2)),floor(locs(ii,1)),2) = 255; % green dot
end

end
```

The MATLAB Function block must have a defined size for the output array. A fast way to define the output size is to copy the input to the output before calling the helper function. This is the code inside the MATLAB Function block:

```
function Y = fcn(I,minContrast)
    coder.extrinsic('FASTHelper');
    Y = I;
    Y = FASTHelper(I,minContrast);
end
```

Implementation for HDL

The FAST algorithm implemented in the Vision HDL Toolbox Corner Detector block in this model tests 9 contiguous pixels from a ring of 16 pixels, and compares their values to the center pixel value. A kernel of 7x7 pixels around each test pixel includes the 16-pixel ring. The diagram shows the center pixel and the ring of 16 pixels around it that is used for the test. The ring pixels, clockwise from the top-middle, are

```
indices = [22 29 37 45 46 47 41 35 28 21 13 5 4 3 9 15];
```

These pixel indices are used for selection and comparison. The order must be contiguous, but the ring can begin at any point.

1	8	15	22	29	36	43
2	9	16	23	30	37	44
3	10	17	24	31	38	45
4	11	18	25	32	39	46
5	12	19	26	33	40	47
6	13	20	27	34	41	48
7	14	21	28	35	42	49

After computing corner metrics using these rings of pixels, the algorithm determines the maximum corner metric in each region and suppresses other detected corners. The model then overlays the non-suppressed corner markers onto the original input image.

The hardware algorithm is in the FASTHDLAlgorithm subsystem. This subsystem supports HDL code generation.

```
open_system([modelname '/FASTHDLAlgorithm'],'force');
```
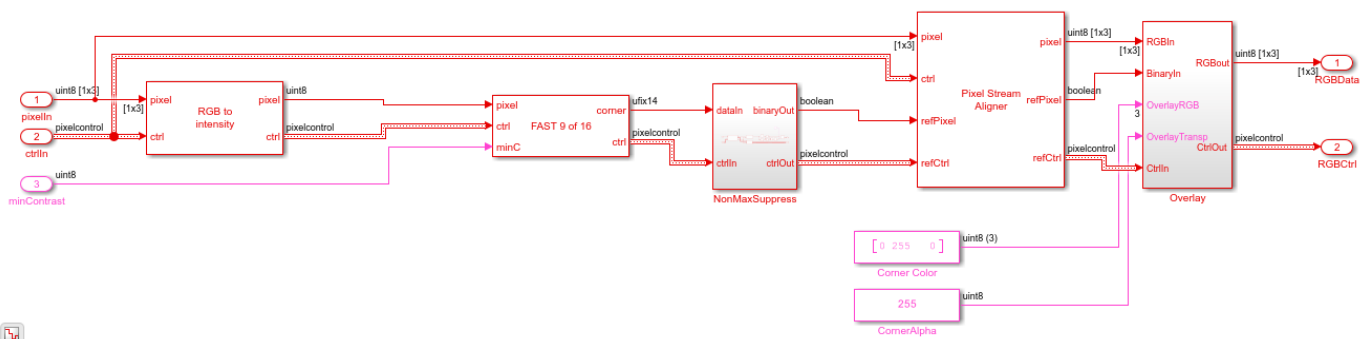

Corner Detection

To determine the presence of a corner, look for all possible 9-pixel contiguous segments of the ring that have values either greater than or less than the threshold value.

In hardware, you can perform all these comparisons in parallel. Each comparator block expands to 16 comparators. The output of the block is 16 binary decisions representing each segment of the ring.

Non-Maximal Suppression

The FAST algorithm identifies many, many potential corners. To reduce subsequent processing, all corners except the corners with the maximum corner metric in a particular region can be removed or suppressed. There are many algorithms for non-maximal suppression suitable for software implementation, but few suitable for hardware. In software, a gradient-based approach is used, which

can be resource intensive in hardware. In this model a simple but very effective technique is to compare corner metrics in a `5x5` kernel and produce a boolean result. The boolean output is true if the corner metric in the center of the kernel is greater than zero (i.e. it is a corner) and also it is the maximum of all the other corner metrics in the `5x5` region. The greater-than-zero condition matches setting `minQuality` to `1` for the `detectFASTFeatures` function.

Since the processing of the pixel stream is from left to right and top to bottom, the results contain some directional effects, such as that the detected corners do not always perfectly align with the objects. The NonMaxSuppress subsystem includes a constant block that allows you to disable suppression and visualize the complete results.

`open_system([modelname '/FASTHDLAlgorithm/NonMaxSuppress'],'force');`
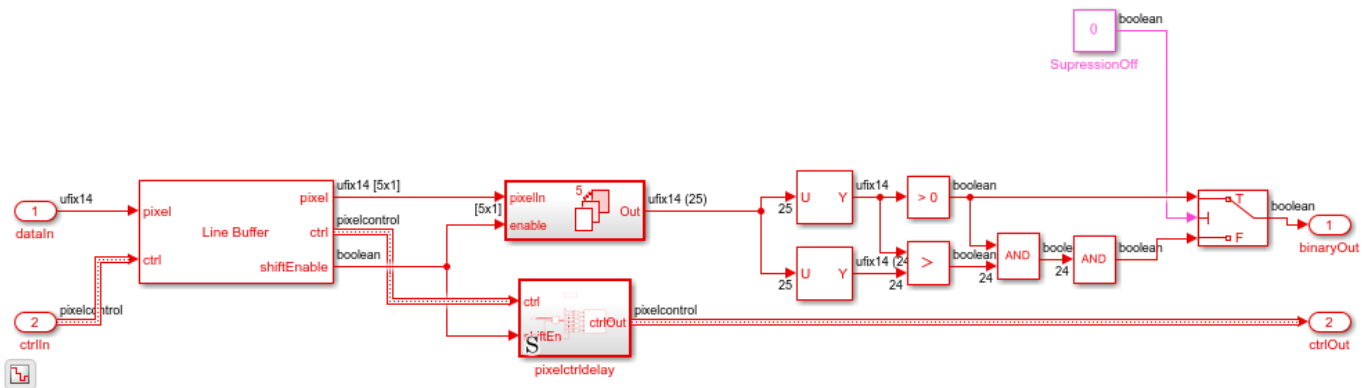
Align and Overlay

At the output of the NonMaxSuppress subsystem, the pixel stream includes markers for the strongest corner in each 5x5 region. Next, the model realigns the detected corners with the original pixel stream using the Pixel Stream Aligner block. After the original stream and the markers are aligned in time, the model overlays a green dot on the corners. The Overlay subsystem contains an alpha mixer with constants for the color and alpha values.

The output viewers show the overlaid green dots for corners detected. The Behavioral Video Viewer shows the output of the `detectFastFeatures` function, and the HDL Video Viewer shows the output of the HDL algorithm.
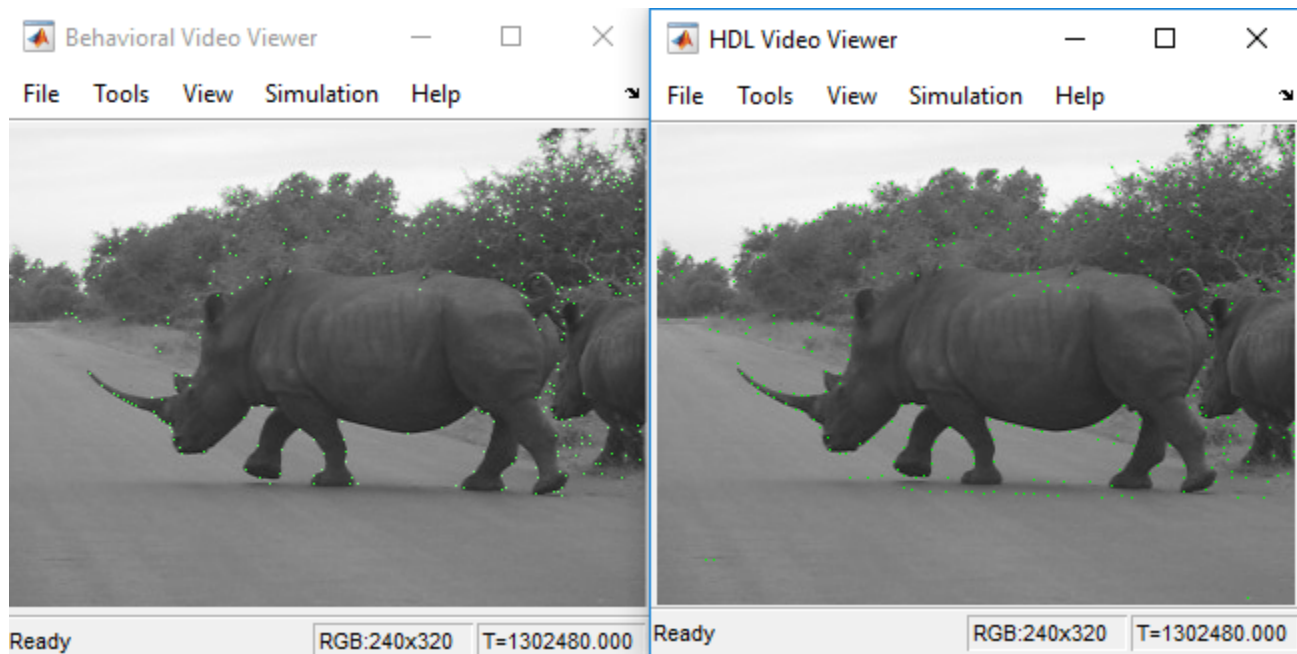
Going Further

The non-maximal suppression algorithm could be improved by following gradients and using a multiple-pass strategy, but that computation would also use more hardware resources.

Conclusion

This example shows how to start using `detectFASTFeatures` in MATLAB and then move to Simulink for the FPGA portion of the design. The hardware algorithm in the Corner Detector block includes a test of the ring around the central pixel in a kernel, and a corner strength metric. The model uses a non-maximal suppression function to remove all but the strongest detected corners. The design then overlays the corner locations onto the original video input, highlighting the corners in green.

References

[1] Rosten, E., and T. Drummond. "Fusing Points and Lines for High Performance Tracking" Proceedings of the IEEE International Conference on Computer Vision, Vol. 2 (October 2005): pp. 1508-1511.

[2] Rosten, E., and T. Drummond. "Machine Learning for High-Speed Corner Detection" Computer Vision - ECCV 2006 Lecture Notes in Computer Science, 2006, 430-43. doi:10.1007/11744023_34.

Lane Detection

This example shows how to implement a lane-marking detection algorithm for FPGAs.

Lane detection is a critical processing stage in Advanced Driving Assistance Systems (ADAS). Automatically detecting lane boundaries from a video stream is computationally challenging and therefore hardware accelerators such as FPGAs and GPUs are often required to achieve real time performance.

In this example model, an FPGA-based lane candidate generator is coupled with a software-based polynomial fitting engine, to determine lane boundaries.

Download Input File

This example uses the visionhdl_caltech.avi file as an input. The file is approximately 19 MB in size. Download the file from the MathWorks website and unzip the downloaded file.

```
laneZipFile = matlab.internal.examples.downloadSupportFile('visionhdl_hdlcoder','caltech_dataset
[outputFolder,~,~] = fileparts(laneZipFile);
unzip(laneZipFile,outputFolder);
caltechVideoFile = fullfile(outputFolder,'caltech_dataset');
addpath(caltechVideoFile);
```

System Overview

The LaneDetectionHDL.slx system is shown below. The HDLLaneDetector subsystem represents the hardware accelerated part of the design, while the SWLaneFitandOverlay subsystem represent the software based polynomial fitting engine. Prior to the Frame to Pixels block, the RGB input is converted to intensity color space.

```
modelname = 'LaneDetectionHDL';
open_system(modelname);
set_param(modelname,'SampleTimeColors','on');
set_param(modelname,'SimulationCommand','Update');
set_param(modelname,'Open','on');
set(allchild(0),'Visible','off');
```
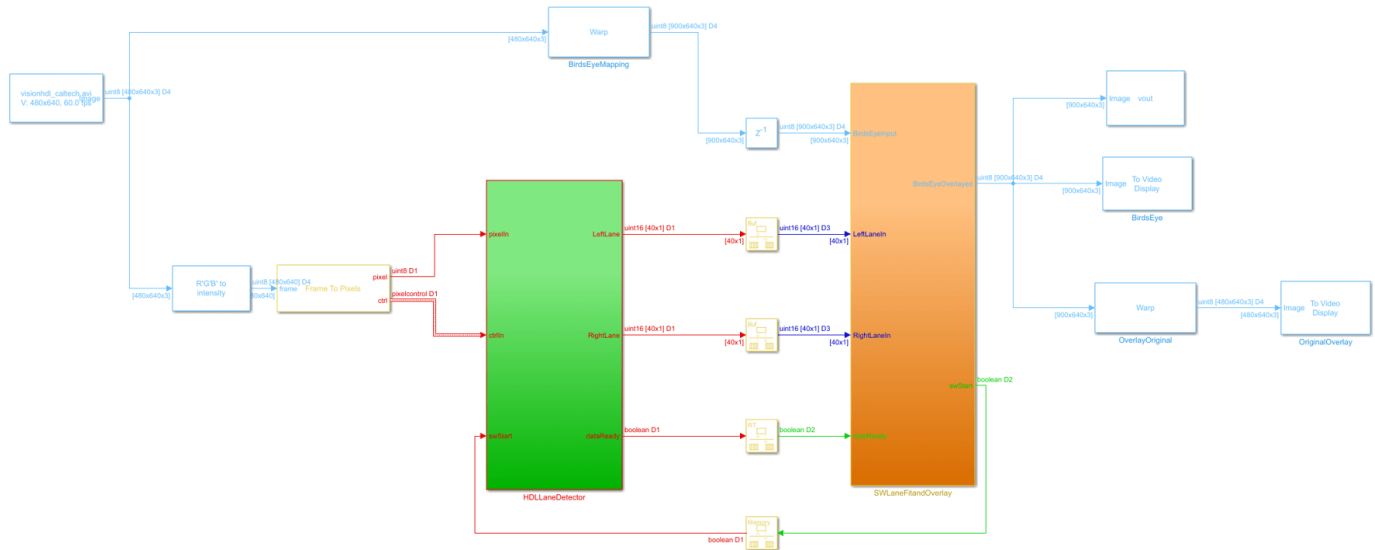
HDL Lane Detector

The HDL Lane Detector represents the hardware-accelerated part of the design. This subsystem receives the input pixel stream from the front-facing camera source, transforms the view to obtain the birds-eye view, locates lane marking candidates from the transformed view and then buffers them up into a vector to send to the software side for curve fitting and overlay.

```
set_param(modelname, 'SampleTimeColors', 'off');
open_system([modelname '/HDLLaneDetector'],'force');
```
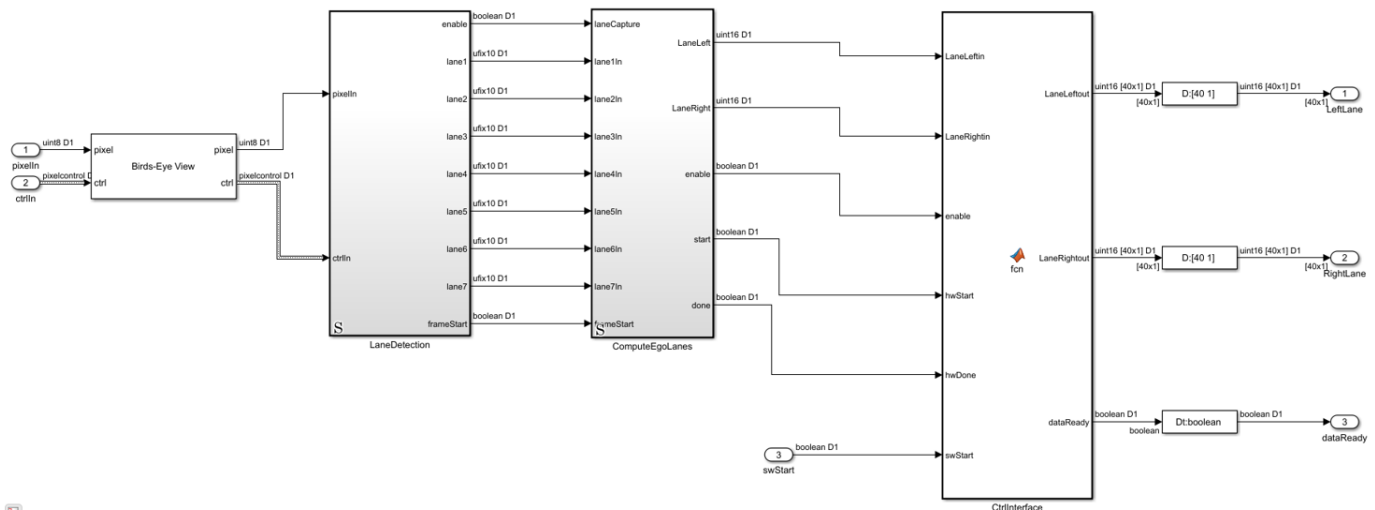
Birds-Eye View

The Birds-Eye View block transforms the front-facing camera view to a birds-eye perspective. Working with the images in this view simplifies the processing requirements of the downstream lane detection algorithms. The front-facing view suffers from perspective distortion, causing the lanes to converge at the vanishing point. The perspective distortion is corrected by applying an inverse perspective transform.

The Inverse Perspective Mapping (IPM) is given by the following expression:

$$(\hat{x}, \hat{y}) = round\left(\frac{h_{11}x + h_{12}y + h_{13}}{h_{31}x + h_{32}y + h_{33}}, \frac{h_{21}x + h_{22}y + h_{23}}{h_{31}x + h_{32}y + h_{33}}\right)$$

The homography matrix, *h*, is derived from four intrinsic parameters of the physical camera setup, namely the focal length, pitch, height, and principle point (from a pinhole camera model). For more details, refer to the Computer Vision Toolbox™ documentation.

You can estimate the homography matrix by using the Computer Vision Toolbox™ `estgeotform2d` function or the Image Processing Toolbox™ `fitgeotform2d` function to create a `projtform2d` object. These functions require a set of matched points between the source frame and birds-eye view frame. The source frame points are taken as the vertices of a trapezoidal region of interest, and can extend past the source frame limits to capture a larger region. For the trapezoid shown the point mapping is:
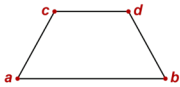
$$sourcePoints = [c_x, c_y;\ d_x, d_y;\ a_x, a_y;\ b_x, b_y]$$

$$birdsEyePoints = [1, 1;\ bAPPL, 1;\ 1, bAVL;\ bAPPL, bAVL]$$

Where *bAPPL* and *bAVL* are the birds-eye view active pixels per line and active video lines respectively.

Direct evaluation of the source (front-facing) to destination (birds-eye) mapping in real time on FPGA/ASIC hardware is challenging. The requirement for division along with the potential for non-sequential memory access from a frame buffer mean that the computational requirements of this part of the design are substantial. Therefore instead of directly evaluating the IPM calculation in real time, an offline analysis of the input to output mapping has been performed and used to pre-compute a mapping scheme. This is possible as the homography matrix is fixed after factory calibration/installation of the camera, due to the camera position, height and pitch being fixed.

In this particular example, the birds-eye output image is a frame of [700x640] dimensions, whereas the front-facing input image is of [480x640] dimensions. There is not sufficient blanking available in order to output the full birds-eye frame before the next front-facing camera input is streamed in. The Birds-Eye view block will therefore not accept any new frame data until it has finished processing the current birds-eye frame.

```
open_system([modelname '/HDLLaneDetector'],'force');
```

Line Buffering and Address Computation

A full sized projective transformation from input to output would result in a [900x640] output image. This requires that the full [480x640] input image is stored in memory, while the source pixel location is calculated using the source location and homography matrix. Ideally on-chip memory should be used for this purpose, removing the requirement for an off-chip frame buffer.

You can determine the number of lines to buffer on-chip by performing inverse row mapping using the homography matrix. The following script calculates the homography matrix from the point mapping, using it to an inverse transform to map the source frame rows to birds-eye view rows.

```
% Source & Birds-Eye Frame Parameters
%   AVL:  Active Video Lines, APPL: Active Pixels Per Line
sAVL  = 480;
sAPPL = 640;
% Birds-Eye Frame
bAVL  = 700;
bAPPL = 640;

% Determine Homography Matrix
%   Point Mapping [NW; NE; SW; SE]
sourcePoints   = [218,196; 421,196; -629,405; 1276,405];
birdsEyePoints = [001,001; 640,001;  001,900;  640,900];
%   Estimate Transform
tf = estgeotform2d(sourcePoints,birdsEyePoints,'projective');
%   Homography Matrix
h = tf.T;

% Visualize Birds-Eye ROI on Source Frame
vidObj  = VideoReader('visionhdl_caltech.avi');
vidFrame = readFrame(vidObj);
vidFrameAnnotated = insertShape(vidFrame,'Polygon',[sourcePoints(1,:) ...
    sourcePoints(2,:) sourcePoints(4,:) sourcePoints(3,:)],          ...
    'LineWidth',5,'Color','red');
vidFrameAnnotated = insertShape(vidFrameAnnotated,'FilledPolygon',   ...
    [sourcePoints(1,:) sourcePoints(2,:) sourcePoints(4,:)           ...
    sourcePoints(3,:)],'LineWidth',5,'Color','red','Opacity',0.2);
figure(1);
subplot(2,1,1);
imshow(vidFrameAnnotated)
title('Source Video Frame');

% Determine Required Birds-Eye Line Buffer Depth
%   Inverse Row Mapping at Frame Centre
x = round(sourcePoints(2,1)-((sourcePoints(2,1)-sourcePoints(1,1))/2));
Y = zeros(1,bAVL);
for ii = 1:1:bAVL
    [~,Y(ii)] = transformPointsInverse(tf,x,ii);
end
numRequiredRows = ceil(Y(0.98*bAVL) - Y(1));

% Visualize Inverse Row Mapping
subplot(2,1,2);
plot(Y,'HandleVisibility','off');   % Inverse Row Mapping
xline(0.98*bAVL,'r','98%','LabelHorizontalAlignment','left',         ...
    'HandleVisibility','off');       % Line Buffer Depth
yline(Y(1),'r--','HandleVisibility','off')
yline(Y(0.98*bAVL),'r')
```

```
title('Birds-Eye View Inverse Row Mapping');
xlabel('Output Row');
ylabel('Input Row');
legend(['Line Buffer Depth: ',num2str(numRequiredRows),' lines'],     ...
    'Location','northwest');
axis equal;
grid on;
```


Source Video Frame

Birds-Eye View Inverse Row Mapping
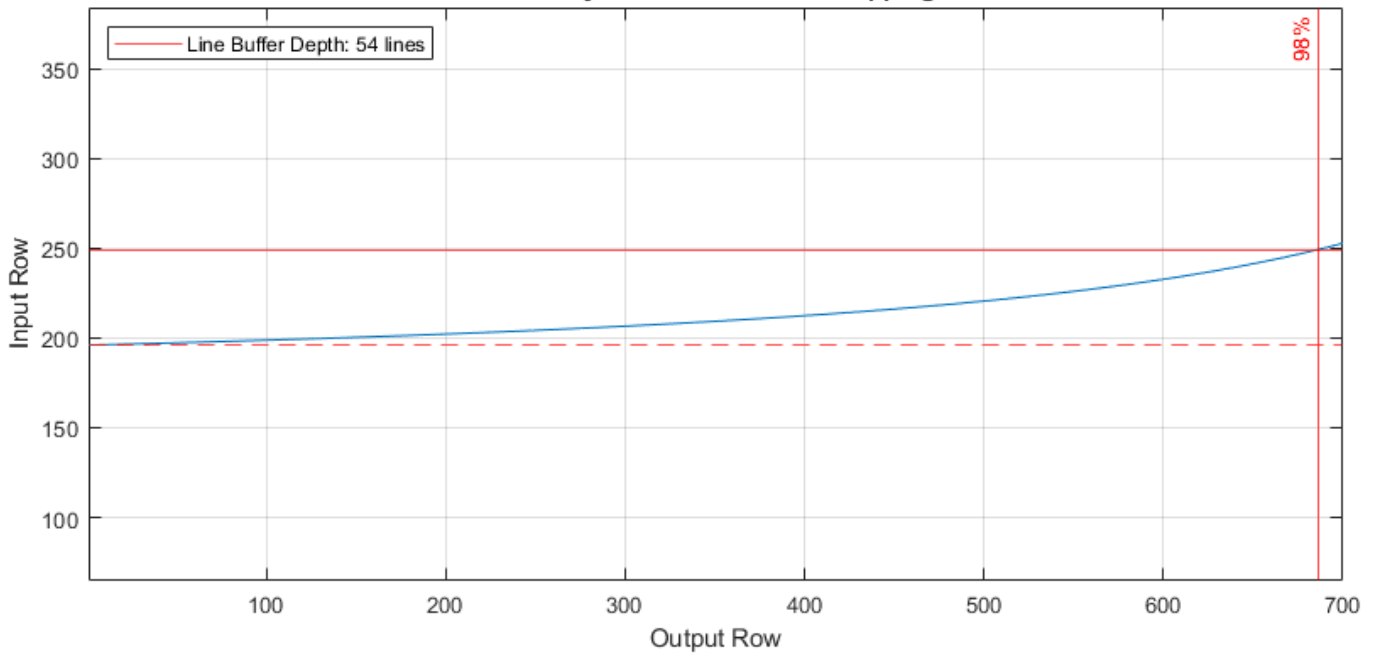
The plot shows the mapping of input line to output line revealing that in order to generate the first 700 lines of the top down birds eye output image, around 50 lines of the input image are required. This is an acceptable number of lines to store using on-chip memory.

Lane Detection

With the birds-eye view image obtained, the actual lane detection can be performed. There are many techniques which can be considered for this purpose. To achieve an implementation which is robust, works well on streaming image data and which can be implemented in FPGA/ASIC hardware at reasonable resource cost, this example uses the approach described in [1]. This algorithm performs a full image convolution with a vertically oriented first order Gaussian derivative filter kernel, followed by sub-region processing.

```
open_system([modelname '/HDLLaneDetector/LaneDetection'],'force');
```
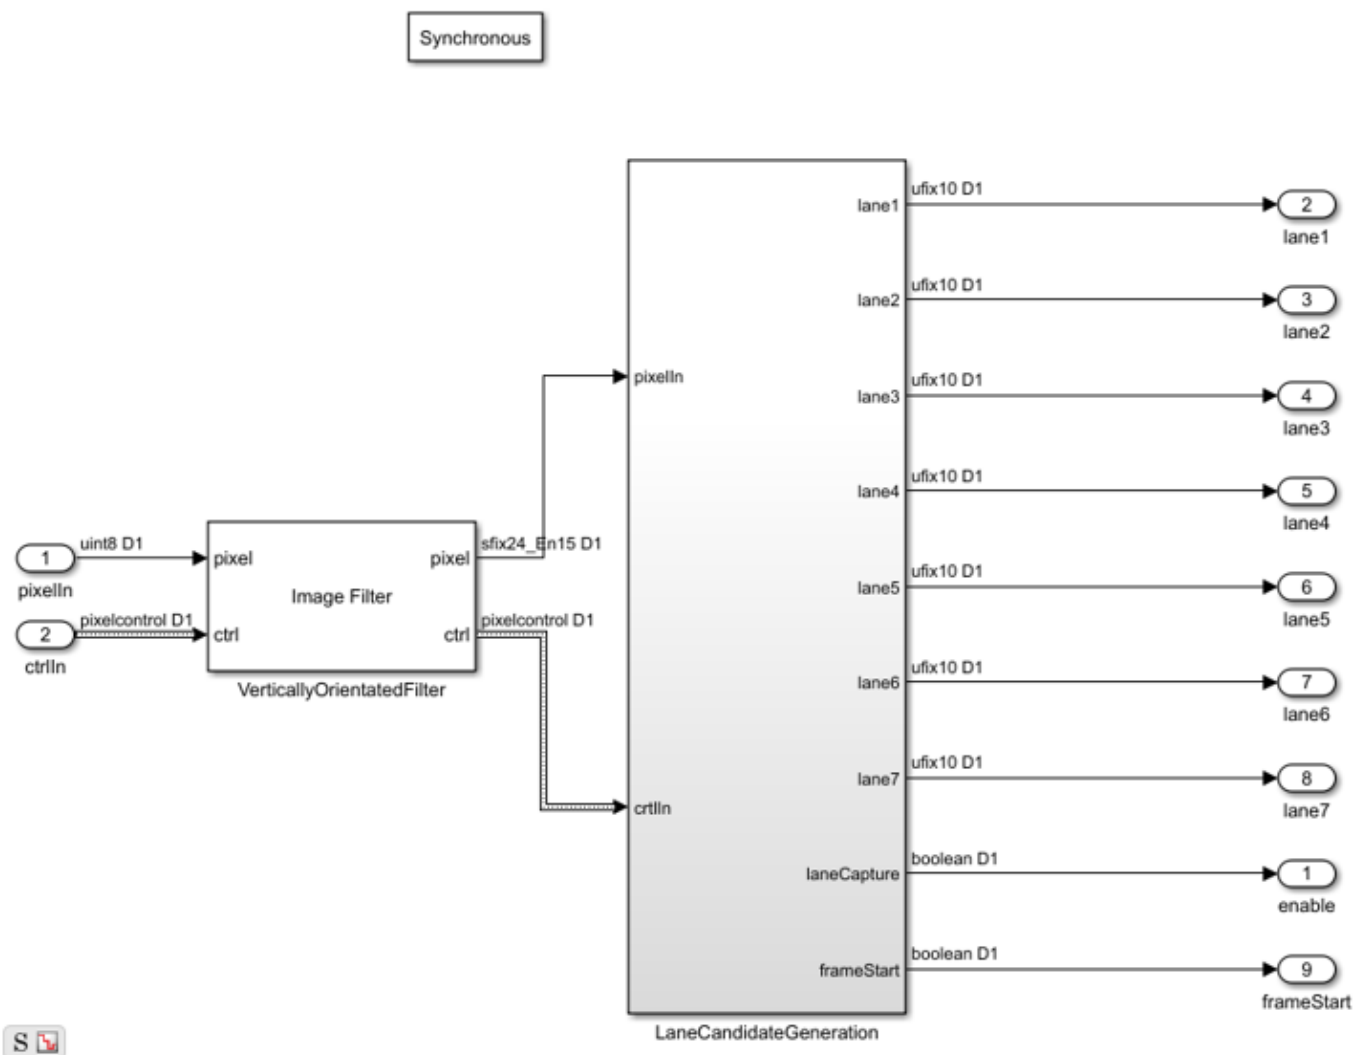

Vertically Oriented Filter Convolution

Immediately following the birds-eye mapping of the input image, the output is convolved with a filter designed to locate strips of high intensity pixels on a dark background. The width of the kernel is 8

pixels, which relates to the width of the lines that appear in the birds-eye image. The height is set to 16 which relates to the size of the dashed lane markings which appear in the image. As the birds-eye image is physically related to the height, pitch etc. of the camera, the width at which lanes appear in this image is intrinsically related to the physical measurement on the road. The width and height of the kernel may need to be updated when operating the lane detection system in different countries.
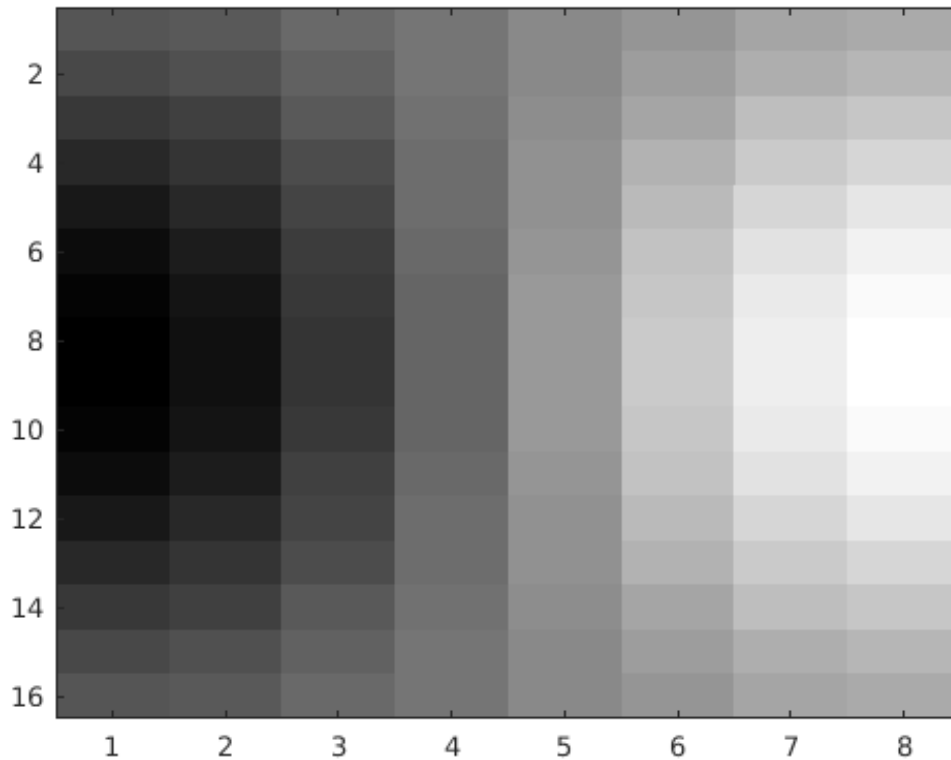
The output of the filter kernel is shown below, using jet colormap to highlight differences in intensity. Because the filter kernel is a general, vertically oriented Gaussian derivative, there is some response from many different regions. However, for the locations where a lane marking is present, there is a strong positive response located next to a strong negative response, which is consistent across columns. This characteristic of the filter output is used in the next stage of the detection algorithm to locate valid lane candidates.
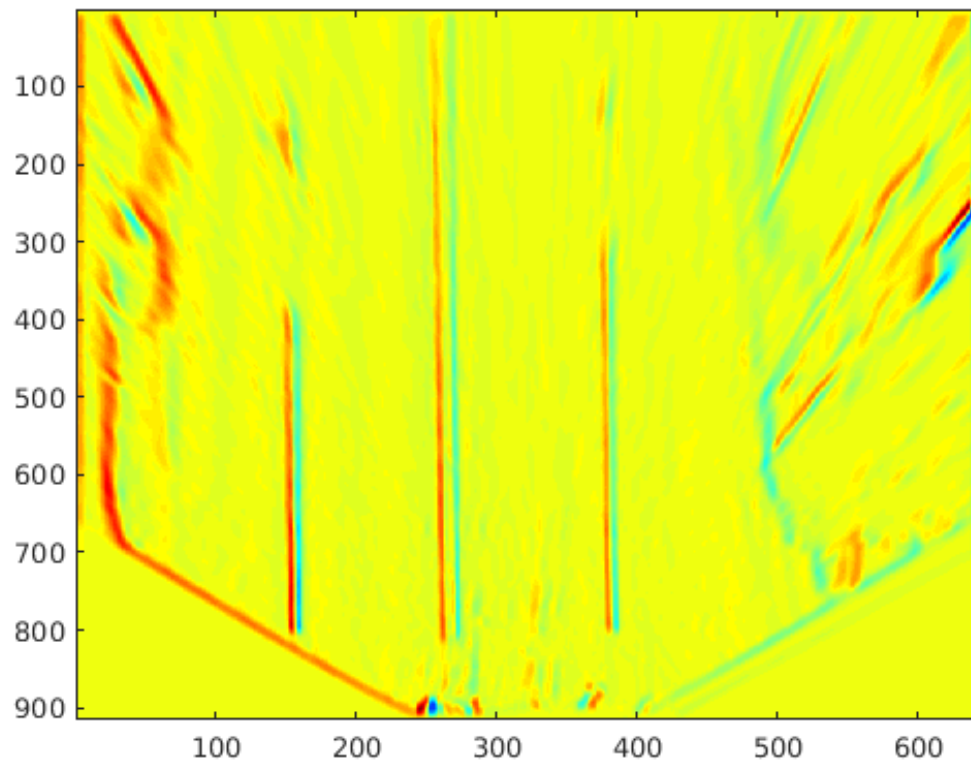
Lane Candidate Generation

After convolution with the Gaussian derivative kernel, sub-region processing of the output is performed in order to find the coordinates where a lane marking is present. Each region consists of 18 lines, with a ping-pong memory scheme in place to ensure that data can be continuously streamed through the subsystem.

```
%Seeing as
open_system([modelname '/HDLLaneDetector/LaneDetection/LaneCandidateGeneration'],'force');
```
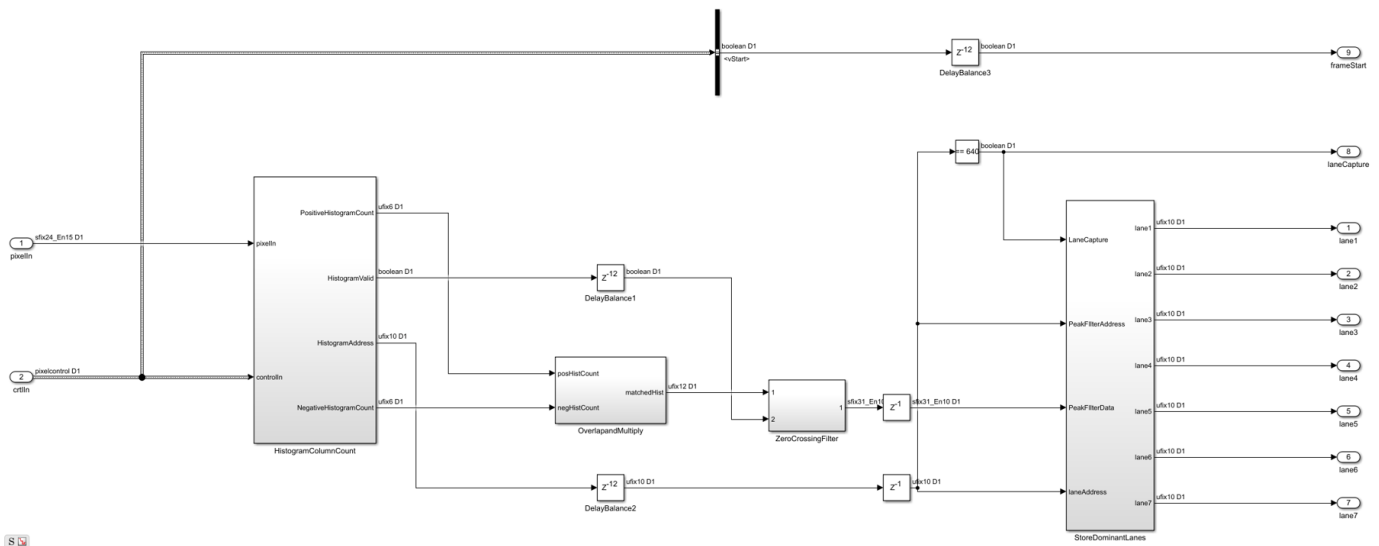
Histogram Column Count

Firstly, HistogramColumnCount counts the number of thresholded pixels in each column over the 18 line region. A high column count indicates that a lane is likely present in the region. This count is performed for both the positive and the negative thresholded images. The positive histogram counts are offset to account for the kernel width. Lane candidates occur where the positive count and negative counts are both high. This exploits the previously noted property of the convolution output where positive tracks appear next to negative tracks.

Internally, the column counting histogram generates the control signalling that selects an 18 line region, computes the column histogram, and outputs the result when ready. A ping-pong buffering scheme is in place which allows one histogram to be reading while the next is writing.

Overlap and Multiply

As noted, when a lane is present in the birds-eye image, the convolution result will produce strips of high-intensity positive output located next to strips of high-intensity negative output. The positive and negative column count histograms locate such regions. In order to amplify these locations, the positive count output is delayed by 8 clock cycles (an intrinsic parameter related to the kernel width), and the positive and negative counts are multiplied together. This amplifies columns where the positive and negative counts are in agreement, and minimizes regions where there is disagreement between the positive and negative counts. The design is pipelined in order to ensure high throughput operation.

Image Histogram - Shifted and Multiplied
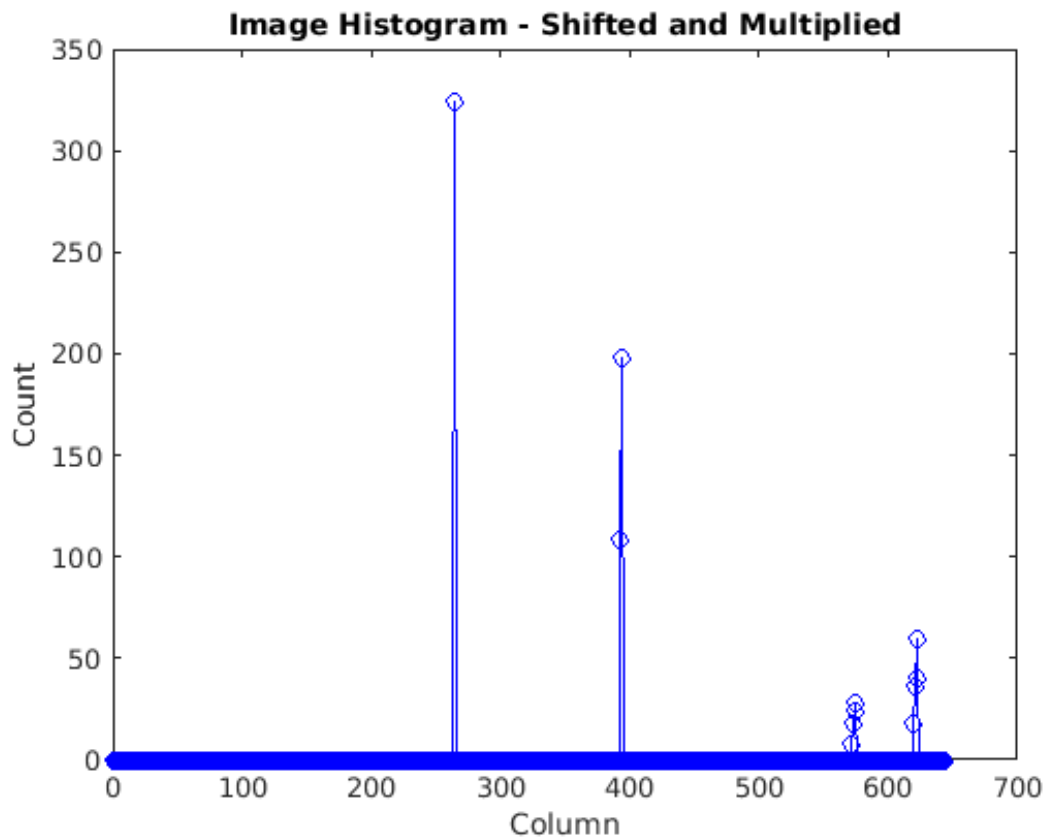
Zero Crossing Filter

At the output of the Overlap and Multiply subsystem, peaks appear where there are lane markings present. A peak detection algorithm determines the columns where lane markings are present. Because the SNR is relatively high in the data, this example uses a simple FIR filtering operation followed by zero crossing detection. The Zero Crossing Filter is implemented using the Discrete FIR Filter block from DSP System Toolbox™. It is pipelined for high-throughput operation.
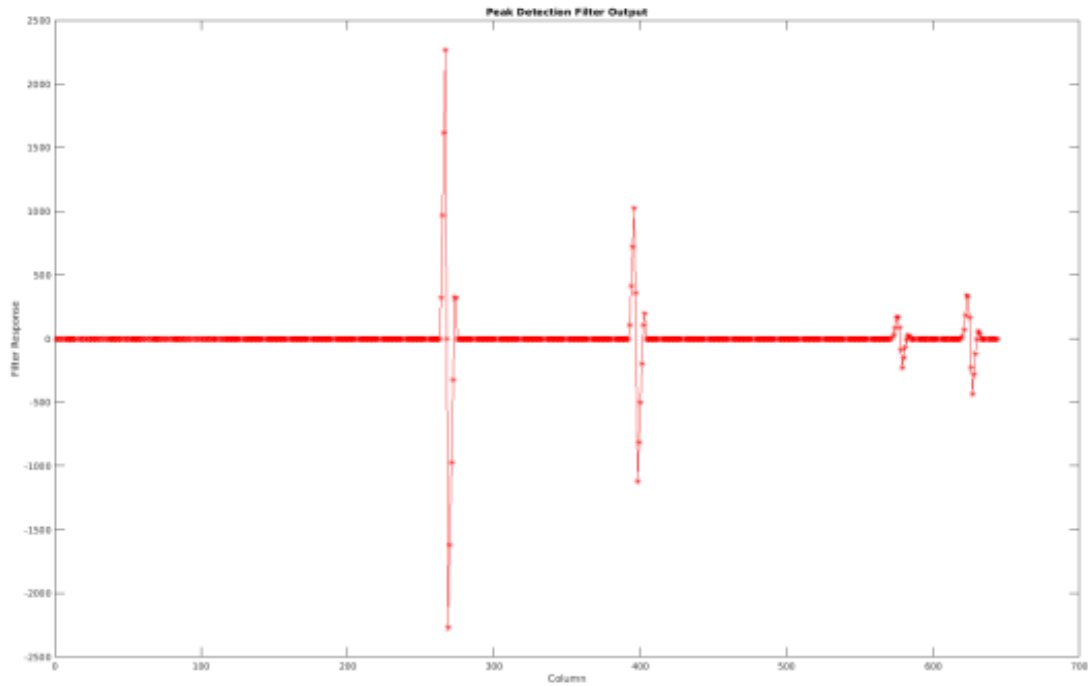
Store Dominant Lanes

The zero crossing filter output is then passed into the Store Dominant Lanes subsystem. This subsystem has a maximum memory of 7 entries, and is reset every time a new batch of 18 lines is reached. Therefore, for each sub-region 7 potential lane candidates are generated. In this subsystem, the Zero Crossing Filter output is streamed through, and examined for potential zero crossings. If a zero crossing does occur, then the difference between the address immediately prior to zero crossing and the address after zero crossing is taken in order to get a measurement of the size of the peak. The subsystem stores the zero crossing locations with the highest magnitude.

```
open_system([modelname '/HDLLaneDetector/LaneDetection/LaneCandidateGeneration/StoreDominantLanes
```
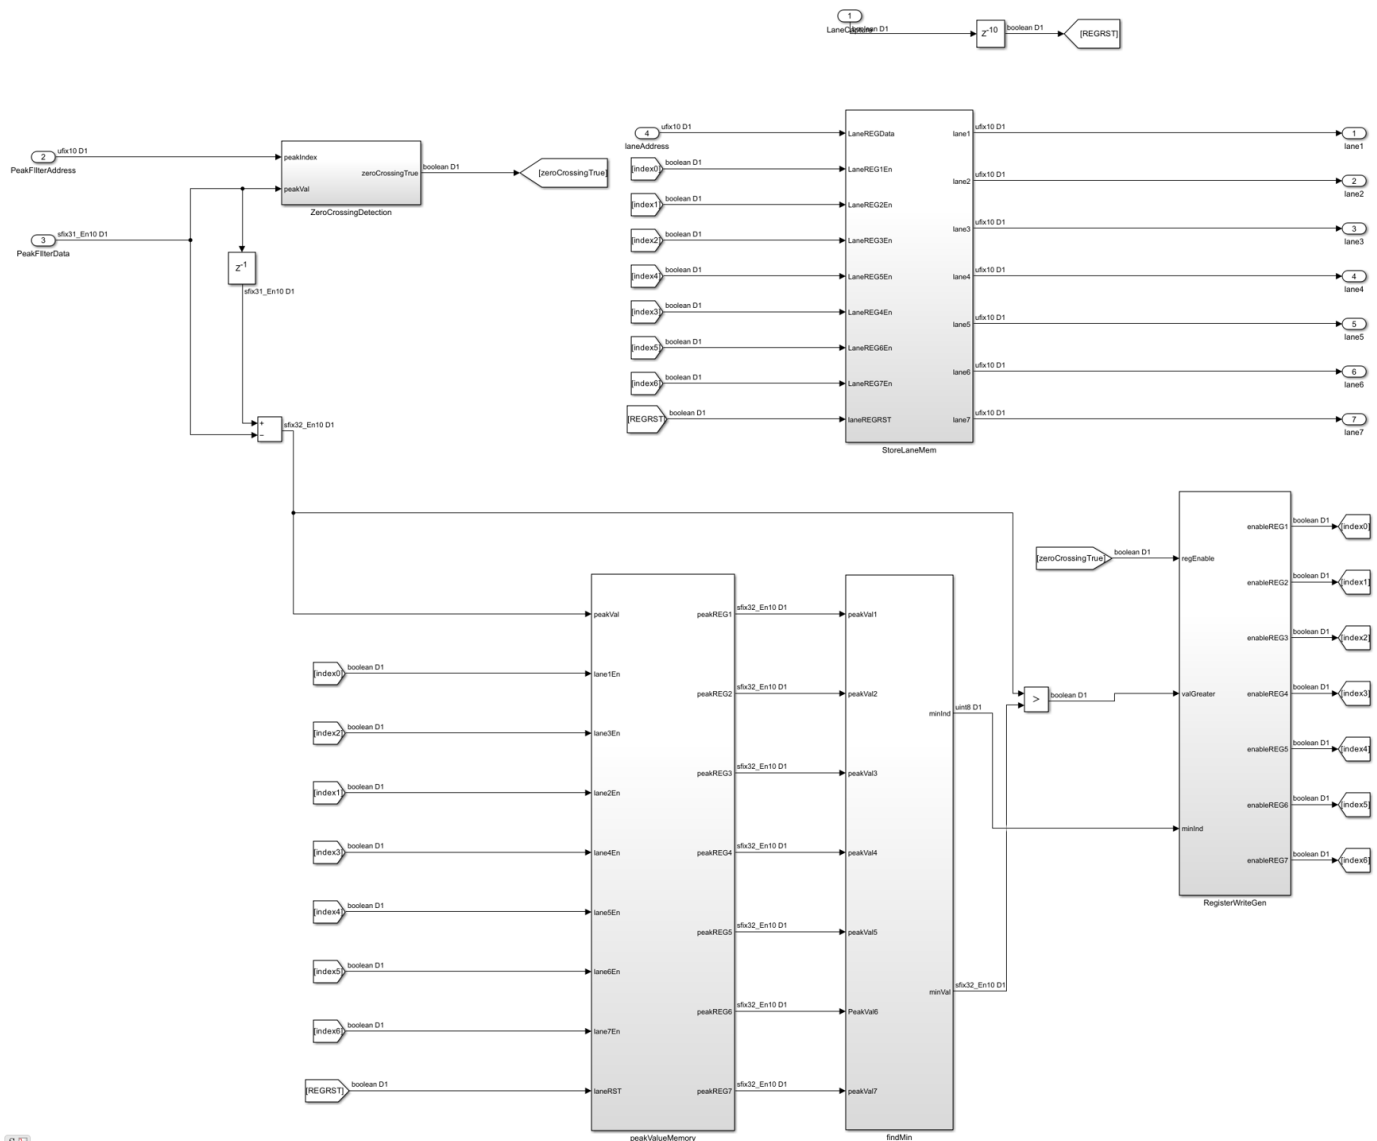
Compute Ego Lanes

The Lane Detection subsystem outputs the 7 most viable lane markings. In many applications, we are most interested in the lane markings that contain the lane in which the vehicle is driving. By computing the so called "Ego-Lanes" on the hardware side of the design, we can reduce the memory bandwidth between hardware and software, by sending 2 lanes rather than 7 to the processor. The Ego-Lane computation is split into two subsystems. The FirstPassEgoLane subsystem assumes that the centre column of the image corresponds to the middle of the lane, when the vehicle is correctly operating within the lane boundaries. The lane candidates which are closest to the center are therefore assumed as the ego lanes. The Outlier Removal subsystem maintains an average width of the distance from lane markings to centre coordinate. Lane markers which are not within tolerance of the current width are rejected. Performing early rejection of lane markers gives better results when performing curve fitting later on in the design.

```
open_system([modelname '/HDLLaneDetector/ComputeEgoLanes'],'force');
```
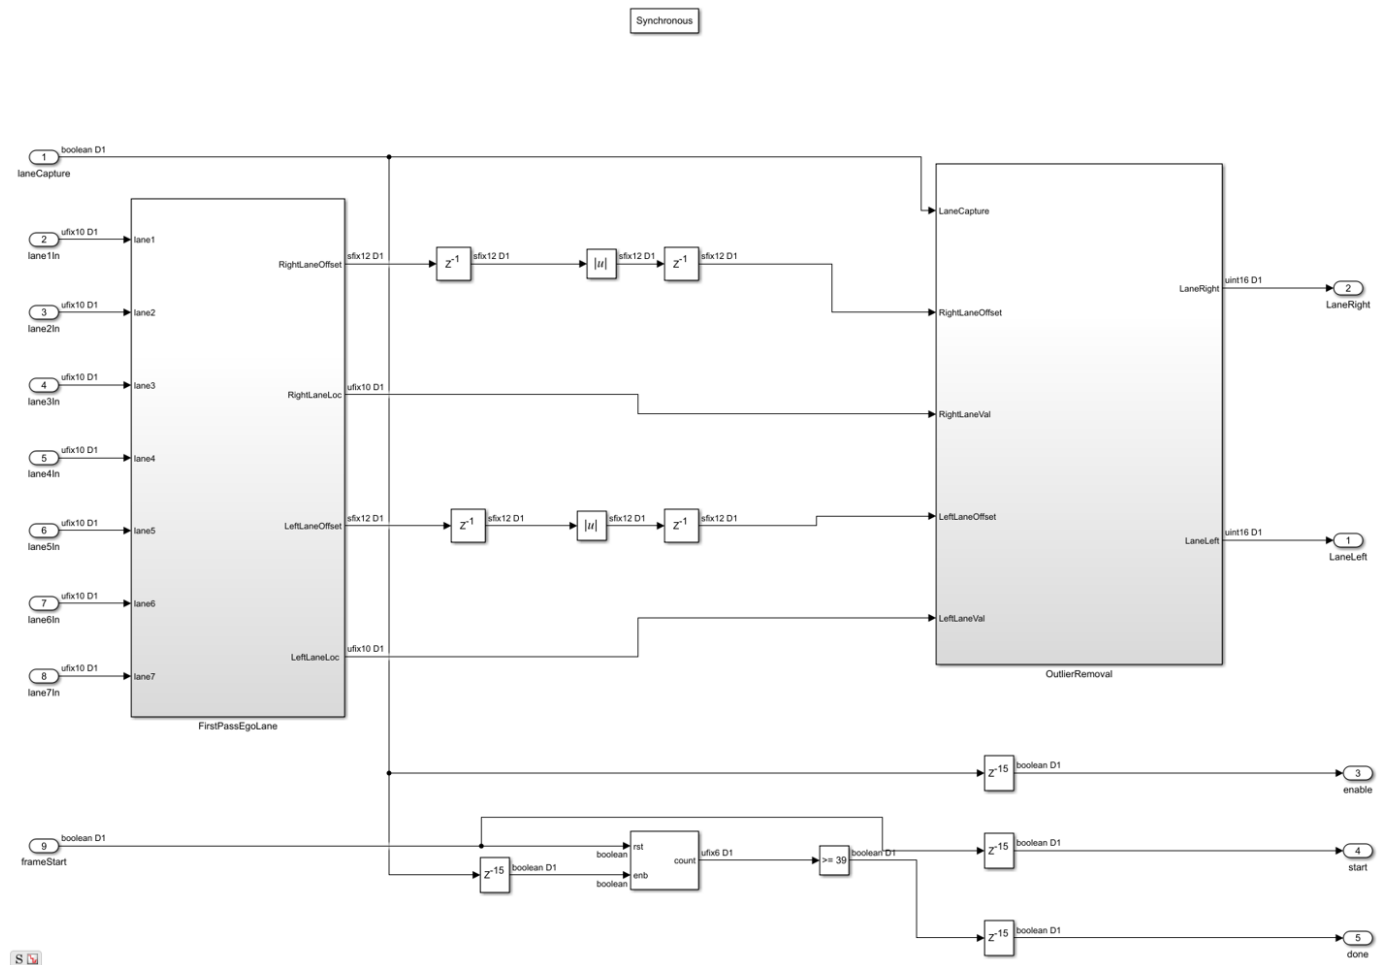
Control Interface

Finally, the computed ego lanes are sent to the CtrlInterface MATLAB function subsystem. This state machine uses the four control signal inputs - enable, hwStart, hwDone, and swStart to determine when to start buffering, accept new lane coordinate into the 40x1 buffer and finally indicate to the software that all 40 lane coordinates have been buffered and so the lane fitting and overlay can be performed. The dataReady signal ensures that software will not attempt lane fitting until all 40 coordinates have been buffered, while the swStart signal ensures that the current set of 40 coordinates will be held until lane fitting is completed.

Software Lane Fit and Overlay

The detected ego-lanes are then passed to the SW Lane Fit and Overlay subsystem, where robust curve fitting and overlay is performed. Recall that the birds-eye output is produced once every two frames or so rather than on every consecutive frame. The curve fitting and overlay is therefore placed in an enabled subsystem, which is only enabled when new ego lanes are produced.

```
open_system([modelname '/SWLaneFitandOverlay'],'force');
```
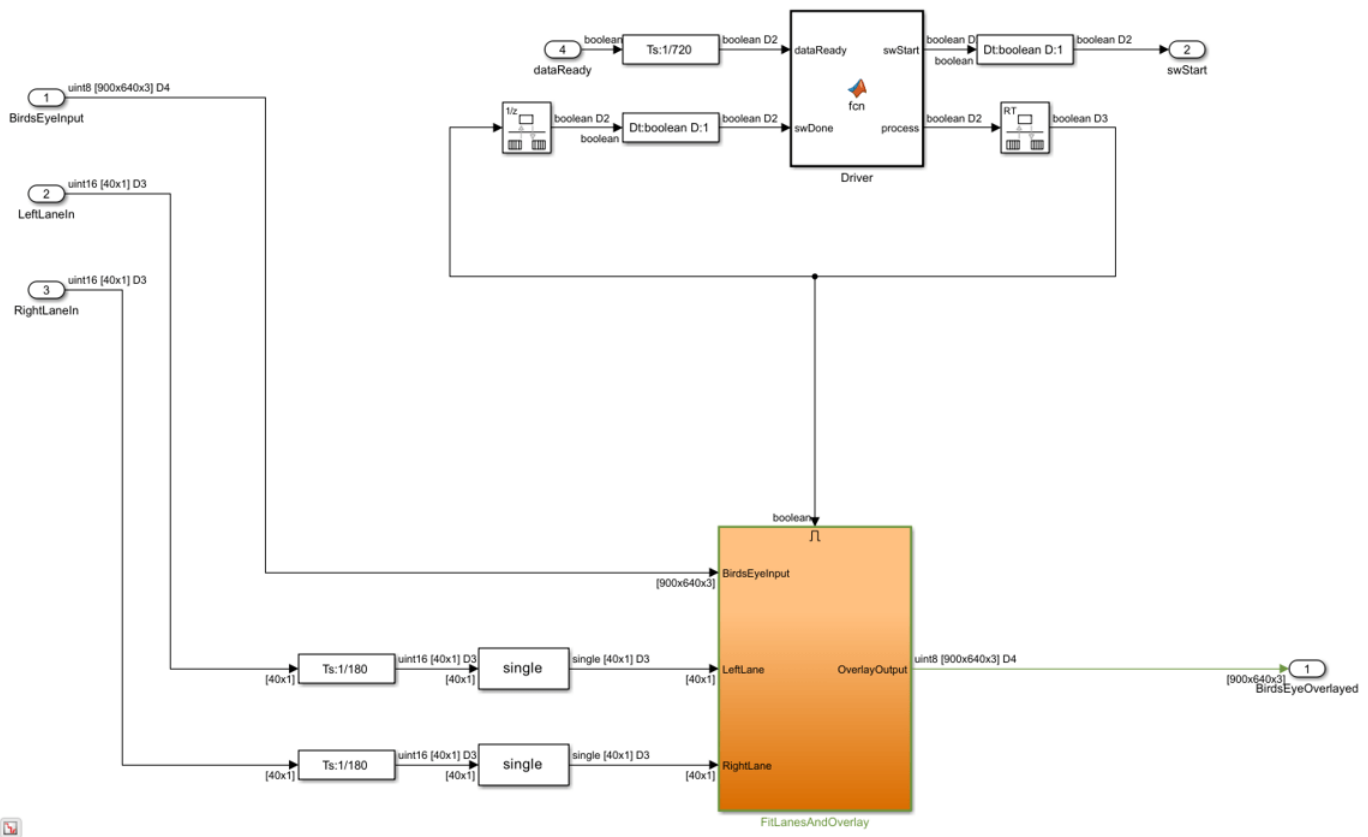
Driver

The Driver MATLAB Function subsystem controls the synchronization between hardware and software. Initially it is in a polling state, where it samples the dataReady input at regular intervals per frame to determine when hardware has buffered a full [40x1] vector of lane coordinates. Once this occurs, it transitions into software processing state where swStart and process outputs are held high. The Driver remains in the software processing state until swDone input is high. Seeing as the process output loops back to swDone input with a rate transition block in between, there is effectively a constant time budget specified for the FitLanesandOverlay subsystem to perform the fitting and overlay. When swDone is high, the Driver will transition into a synchronization state, where swStart is held low to indicate to hardware that processing is complete. The synchronization between software and hardware is such that hardware will hold the [40x1] vector of lane coordinates until the swStart signal transitions back to low. When this occurs, dataReady output of hardware will then transition back to low.

Fit Lanes and Overlay

The Fit Lanes and Overlay subsystem is enabled by the Driver. It performs the necessary arithmetic required in order to fit a polynomial onto the lane coordinate data received at input, and then draws the fitted lane and lane coordinates onto the Birds-Eye image.

Fit Lanes

The Fit Lanes subsystem runs a RANSAC based line-fitting routine on the generated lane candidates. RANSAC is an iterative algorithm which builds up a table of inliers based on a distance measure

between the proposed curve, and the input data. At the output of this subsystem, there is a [3x1] vector which specifies the polynomial coefficients found by the RANSAC routine.

Overlay Lane Markings

The Overlay Lane Markings subsystem performs image visualization operations. It overlays the ego lanes and curves found by the lane-fitting routine.

```
open_system([modelname '/SWLaneFitandOverlay/FitLanesAndOverlay'],'force');
```
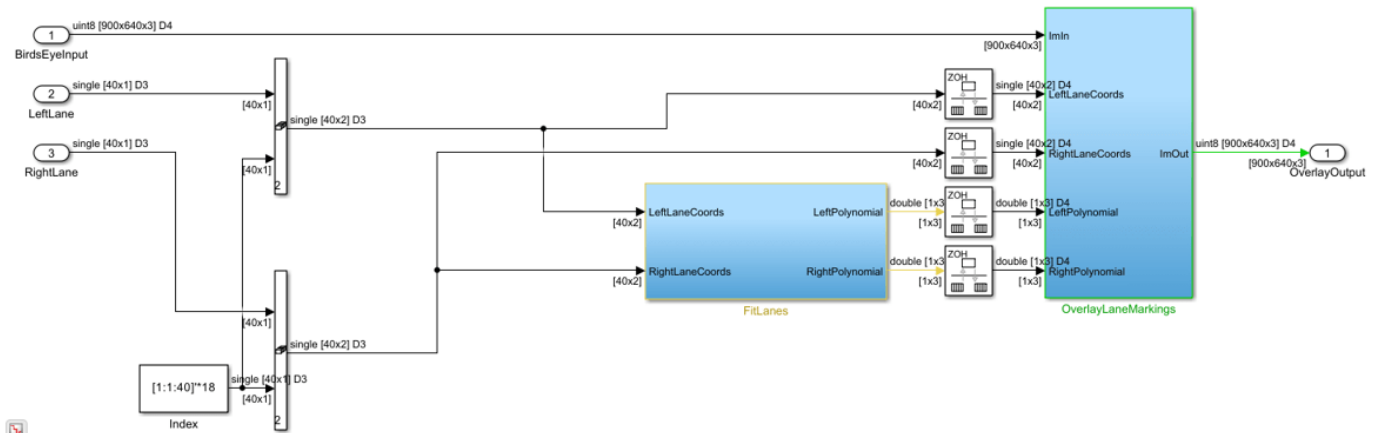

Results of the Simulation

The model includes two video displays shown at the output of the simulation results. The **BirdsEye** display shows the output in the warped perspective after lane candidates have been overlaid, polynomial fitting has been performed and the resulting polynomial overlaid onto the image. The **OriginalOverlay** display shows the **BirdsEye** output warped back into the original perspective.

Due to the large frame sizes used in this model, simulation can take a relatively long time to complete. If you have an HDL Verifier™ license, you can accelerate simulation speed by directly running the HDL Lane Detector subsystem in hardware using FPGA in the Loop.
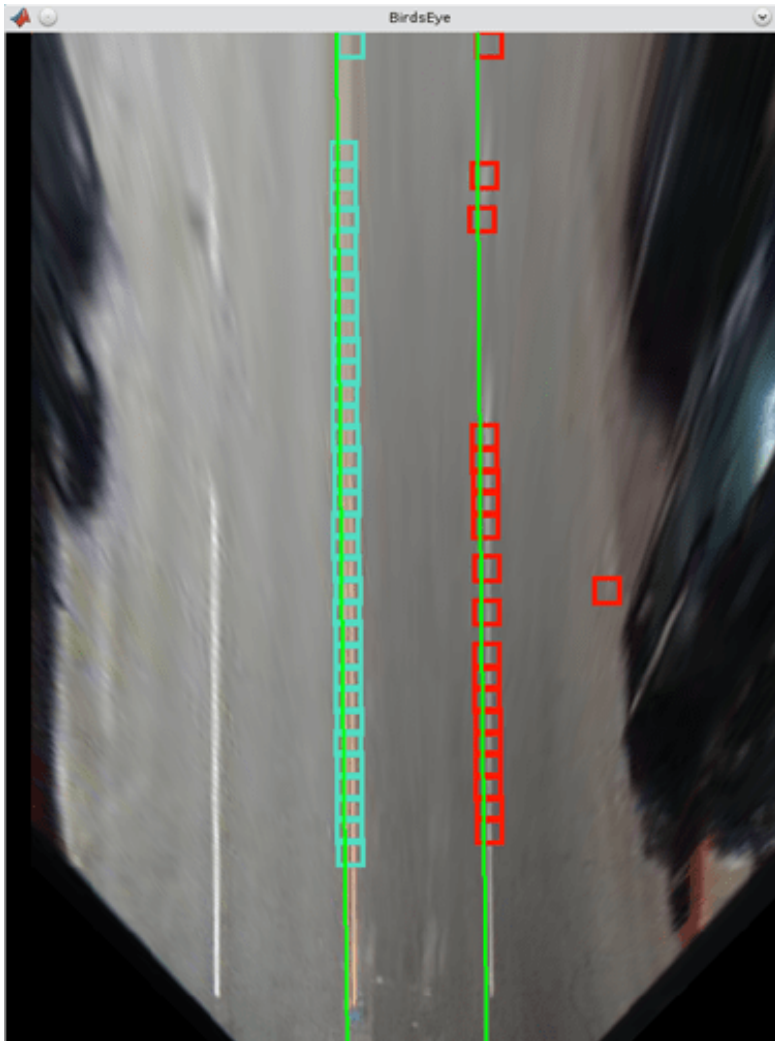
HDL Code Generation

To check and generate the HDL code referenced in this example, you must have an HDL Coder™ license.

To generate the HDL code, use the following command.

```
makehdl('LaneDetectionHDL/HDLLaneDetector')
```

To generate the test bench, use the following command. Note that test bench generation takes a long time due to the large data size. You may want to reduce the simulation time before generating the test bench.

```
makehdltb('LaneDetectionHDL/HDLLaneDetector')
```

For faster test bench simulation, you can generate a SystemVerilog DPIC test bench using the following command.

```
makehdltb('LaneDetectionHDL/HDLLaneDetector','GenerateSVDPITestBench','ModelSim')
```

Conclusion

This example has provided insight into the challenges of designing ADAS systems in general, with particular emphasis paid to the acceleration of critical parts of the design in hardware.

References

[1] R. K. Satzoda and Mohan M. Trivedi, "Vision based Lane Analysis: Exploration of Issues and Approaches for Embedded Realization", 2013 IEEE Conference on Computer Vision and Pattern Recognition.

[2] Video from Caltech Lanes Dataset - Mohamed Aly, "Real time Detection of Lane Markers in Urban Streets", 2008 IEEE Intelligent Vehicles Symposium - used with permission.

See Also

More About

- "Hardware-Software Co-Design Workflow for SoC Platforms" (HDL Coder)

Generate Cartoon Images Using Bilateral Filtering

This example shows how to generate cartoon lines and overlay them onto an image.

Bilateral filtering [1] is used in computer vision systems to filter images while preserving edges and has become ubiquitous in image processing applications. Those applications include denoising while preserving edges, texture and illumination separation for segmentation, and cartooning or image abstraction to enhance edges in a quantized color-reduced image.

Bilateral filtering is simple in concept: each pixel at the center of a neighborhood is replaced by the average of its neighbors. The average is computed using a weighted set of coefficients. The weights are determined by the spatial location in the neighborhood (as in a traditional Gaussian blur filter), and the intensity difference from the center value of the neighborhood.

These two weighting factors are independently controllable by the two standard deviation parameters of the bilateral filter. When the intensity standard deviation is large, the bilateral filter acts more like a Gaussian blur filter, because the intensity Gaussian is less peaked. Conversely, when the intensity standard deviation is smaller, edges in the intensity are preserved or enhanced.

This example model provides a hardware-compatible algorithm. You can generate HDL code from this algorithm, and implement it on a board using a Xilinx™ Zynq™ reference design. See "Bilateral Filtering with Zynq-Based Hardware" (Vision HDL Toolbox Support Package for Xilinx Zynq-Based Hardware).

Download Input File

This example uses the potholes2.avi file as an input. The file is approximately 50 MB in size. Download the file from the MathWorks website and unzip the downloaded file.

```
potholeZipFile = matlab.internal.examples.downloadSupportFile('visionhdl','potholes2.zip');
[outputFolder,~,~] = fileparts(potholeZipFile);
unzip(potholeZipFile,outputFolder);
potholeVideoFile = fullfile(outputFolder,'potholes2');
addpath(potholeVideoFile);
```

Introduction

The BilateralFilterHDLExample.slx system is shown here.

```
modelname = 'BilateralFilterHDLExample';
open_system(modelname);
set_param(modelname, 'SampleTimeColors', 'on');
set_param(modelname,'SimulationCommand','Update');
set_param(modelname, 'Open', 'on');
set(allchild(0),'Visible', 'off');
```
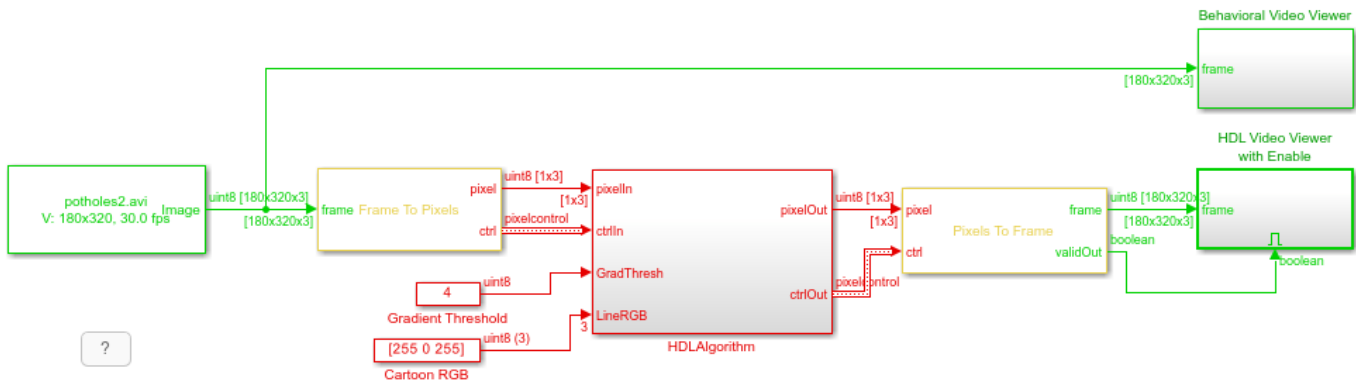
Copyright 2017-2020 The MathWorks, Inc.

Step 1: Establish the Parameter Values

To achieve a modest Gaussian blur of the input, choose a relatively large spatial standard deviation of 3. To give strong emphasis to the edges of the image, choose an intensity standard deviation of 0.75. The intensity Gaussian is built from the image data in the neighborhood, so this plot represents the maximum possible values. Note the small vertical scale on the spatial Gaussian plot.

```
figure('units','normalized','outerposition',[0 0.5 0.75 0.45]);
subplot(1,2,1);
s1 = surf(fspecial('gaussian',[9 9 ],3));
subplot(1,2,2);
s2 = surf(fspecial('gaussian',[9 9 ],0.75));
legend(s1,'Spatial Gaussian 3.0');
legend(s2,'Intensity Gaussian 0.75');
```


Fixed-Point Settings

For HDL code generation, you must choose a fixed-point data type for the filter coefficients. The coefficient type should be an unsigned type. For bilateral filtering, the input range is always assumed to be on the interval $[0, 1]$. Therefore, a `uint8` input with a range of values from $[0, 255]$ are treated as $\frac{[0, 255]}{255}$. The calculated coefficient values are less than 1. The exact values of the coefficients depend on the neighborhood size and the standard deviations. Larger neighborhoods spread the Gaussian function such that each coefficient value is smaller. A larger standard deviation flattens the Gaussian to produce more uniform values, while a smaller standard deviation produces a peaked response.

If you try a type and the coefficients are quantized such that more than half of the kernel becomes zero for all input, the `Bilateral Filter` block issues a warning. If all of the coefficients are zero after quantization, the block issues an error.

Step 2: Filter the Intensity Image

The model converts the incoming RGB image to intensity using the Color Space Converter block. Then the grayscale intensity image is sent to the Bilateral Filter block, which is configured for a 9-by-9 neighborhood and the parameters established previously.

The bilateral filter provides some Gaussian blur but will strongly emphasize larger edges in the image based on the 9-by-9 neighborhood size.

```
open_system([modelname '/HDLAlgorithm'],'force');
```
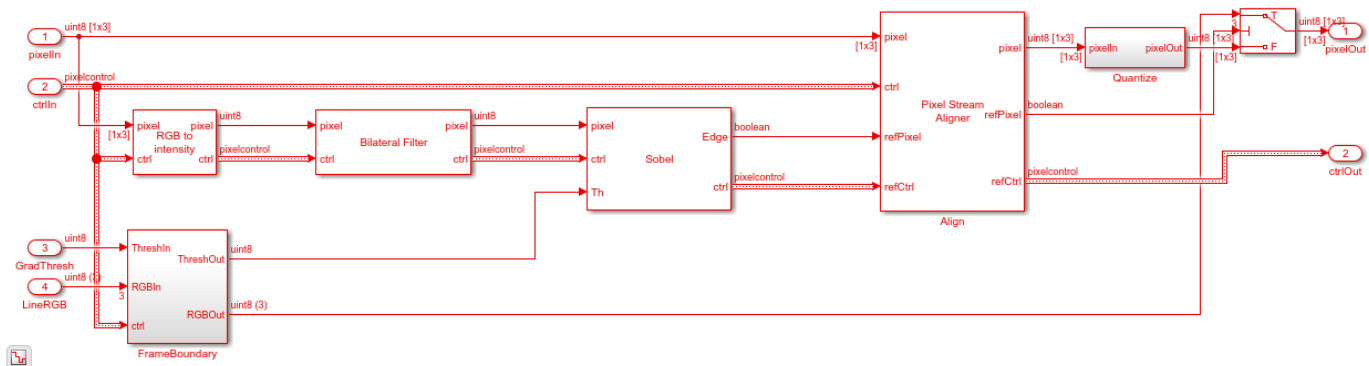

Step 3: Compute Gradient Magnitude

Next, the Sobel Edge Detector block computes the gradient magnitude. Since the image was pre-filtered using a bilateral filter with a fairly large neighborhood, the smaller, less important edges in the image will not be emphasized during edge detection.

The threshold parameter for the Sobel Edge Detector block can come from a constant value on the block mask or from a port. The block in this model uses port to allow the threshold to be set dynamically. This threshold value must be computed for your final system, but for now, you can just choose a good value by observing results.

Synchronize the Computed Edges

To overlay the thresholded edges onto the original RGB image, you must realign the two streams. The processing delay of the bilateral filter and edge detector means that the thresholded edge stream and the input RGB pixel stream are not aligned in time.

The Pixel Stream Aligner block brings them back together. The RGB pixel stream is connected to the upper pixel input port, and the binary threshold image pixel is connected to the reference input port. The block delays the RGB pixel stream to match the threshold stream.

You must set the number of lines parameter to a value that allows for the delay of both the bilateral filter and the edge detector. The 9-by-9 bilateral filter has a delay of more than 4 lines, while the edge detector has a delay of a bit more than 1 line. For safety, set the `Maximum number of lines` to 10 for now so that you can try different neighborhood sizes later. Once your design is done, you can determine the actual number of lines of delay by observing the control signal waveforms.

Color Quantization

Color quantization reduces the number of colors in an image to make processing it easier. Color quantization is primarily a clustering problem, because you want to find a single representative color for a cluster of colors in the original image.

For this problem, you can apply many different clustering algorithms, such as k-means or the median cut algorithm. Another common approach is using octrees, which recursively divide the color space into 8 octants. Normally you set a maximum depth of the tree, which controls the recursive subtrees that will be eliminated and therefore represented by one node in the subtree above.

These algorithms require that you know beforehand all of the colors in the original image. In pixel streaming video, the color discovery step introduces an undesirable frame delay. Color quantization is also generally best done in a perceptually uniform color space such as L*a*b. When you cluster colors in RGB space, there is no guarantee that the result will look representative to a human viewer.

The `Quantize` subsystem in this model uses a much simpler form of color quantization based on the most significant 4 bits of each 8-bit color component. RGB triples with 8-bit components can represent up to $2^{24} = 2^8 \cdot 2^8 \cdot 2^8$ colors but no single image can use all those colors. Similarly when you reduce the number of bits per color to 4, the image can contain up to $2^{12} = 2^4 \cdot 2^4 \cdot 2^4$ colors. In practice a 4-bit-per-color image typically contains only several hundred unique colors.

After shifting each color component to the right by 4 bits, the model shifts the result back to the left by 4 bits to maintain the 24-bit RGB format supported by the video viewer. In an HDL system, the next processing steps would pass on only the 4-bit color RGB triples.

```
open_system([modelname '/HDLAlgorithm/Quantize'],'force');
```
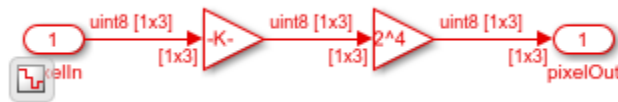

Overlay the Edges

A switch block overlays the edges on the original image by selecting either the RGB stream or an RGB parameter. The switch is flipped based on the edge-detected binary image. Because cartooning requires strong edges, the model does not use an alpha mixer.

Parameter Synchronization

In addition to the pixel and control signals, two parameters enter the HDLAlgorithm subsystem: the gradient threshold and the line RGB triple for the overlay color. The `FrameBoundary` subsystem provides run-time control of the threshold and the line color. However, to avoid an output frame with a mix of colors or thresholds, the subsystem registers the parameters only at the start of each frame.

```
open_system([modelname '/HDLAlgorithm/FrameBoundary'],'force');
```
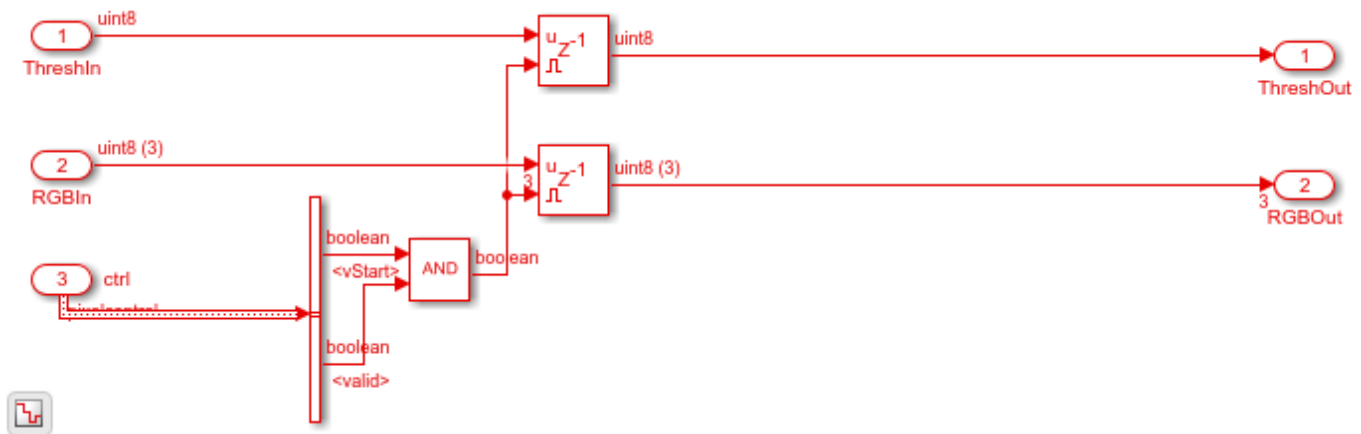
Simulation Results

After you run the simulation, you can see that the resulting images from the simulation show bold lines around the detected features in the input video.

HDL Code Generation

To check and generate the HDL code referenced in this example, you must have an HDL Coder™ license.

To generate the HDL code, use the following command.

```
makehdl('BilateralHDLExample/HDLAlgorithm')
```

To generate the test bench, use the following command. Note that test bench generation takes a long time due to the large data size. Consider reducing the simulation time before generating the test bench.

```
makehdltb('BilateralHDLExample/HDLAlgorithm')
```

The part of the model between the Frame to Pixels and Pixels to Frame blocks can be implemented on an FPGA. The HDLAlgorithm subsystem includes all elements of the bilateral filter, edge detection, and overlay.

Going Further

The bilateral filter in this example is configured to emphasize larger edges while blurring smaller ones. To see the edge detection and overlay without bilateral filtering, right-click the Bilateral Filter block and select `Comment Through`. Then rerun the simulation. The updated results show that many smaller edges are detected and in general, the edges are much noisier.

This model has many parameters you can control, such as the bilateral filter standard deviations, the neighborhood size, and the threshold value. The neighborhood size controls the minimum width of emphasized edges. A smaller neighborhood results in more small edges being highlighted.

You can also control how the output looks by changing the RGB overlay color and the color quantization. Changing the edge detection threshold controls the strength of edges that are overlaid.

To further cartoon the image, you can try adding multiple bilateral filters. With a the right parameters, you can generate a very abstract image that is suitable for a variety of image segmentation algorithms.

Conclusion

This model generated a cartoon image using bilateral filtering and gradient generation. The model overlaid the cartoon lines on a version of the original RGB image that was quantized to a reduced number of colors. This algorithm is suitable for FPGA implementation.

References

[1] Tomasi, C., and R. Manducji. "Bilateral filtering for gray and color images." Sixth International Conference on Computer Vision, 1998.

Pothole Detection

This example extends the "Generate Cartoon Images Using Bilateral Filtering" on page 2-78 example to include calculating a centroid and overlaying a centroid marker and text label on detected potholes.

Road hazard or pothole detection is an important part of any automated driving system. Previous work [1] on automated pothole detection defined a pothole as an elliptical area in the road surface that has a darker brightness level and different texture than the surrounding road surface. Detecting potholes using image processing then becomes the task of finding regions in the image of the road surface that fit the chosen criterion. You can use any or all of the elliptical shape, darker brightness or texture criterion.

To measure the elliptical shape you can use a voting algorithm such as Hough circle, or a template matching algorithm, or linear algebra-based methods such as a least squares fit. Measuring the brightness level is simple in image processing by selecting a brightness segmentation value. The texture can be assessed by calculating the spatial frequency in a region using techniques such as the FFT.

This example uses brightness segmentation with an area metric so that smaller defects are not detected. To find the center of the defect, this design calculates the centroid. The model overlays a marker on the center of the defect and overlays a text label on the image.

Download Input File

This example uses the potholes2.avi file as an input. The file is approximately 50 MB in size. Download the file from the MathWorks website and unzip the downloaded file.

```
potholeZipFile = matlab.internal.examples.downloadSupportFile('visionhdl','potholes2.zip');
[outputFolder,~,~] = fileparts(potholeZipFile);
unzip(potholeZipFile,outputFolder);
potholeVideoFile = fullfile(outputFolder,'potholes2');
addpath(potholeVideoFile);
```

Introduction

The PotHoleHDLDetector.slx system is shown below. The PotHoleHDL subsystem contains the pothole detector and overlay algorithms and supports HDL code generation. There are four input parameters that control the algorithm. The ProcessorBehavioral subsystem writes character maps into a RAM for use as overlay labels.

```
modelname = 'PotHoleHDLDetector';
open_system(modelname);
set_param(modelname,'SampleTimeColors','on');
set_param(modelname,'SimulationCommand','Update');
set_param(modelname,'Open','on');
set(allchild(0),'Visible','off');
```
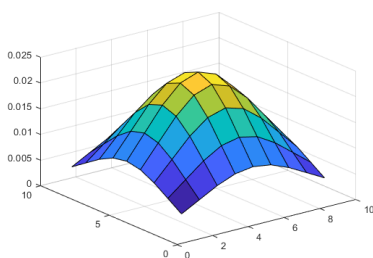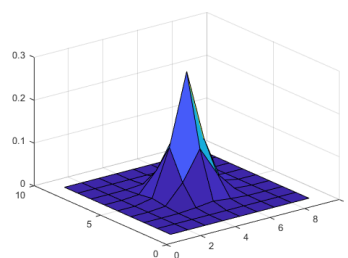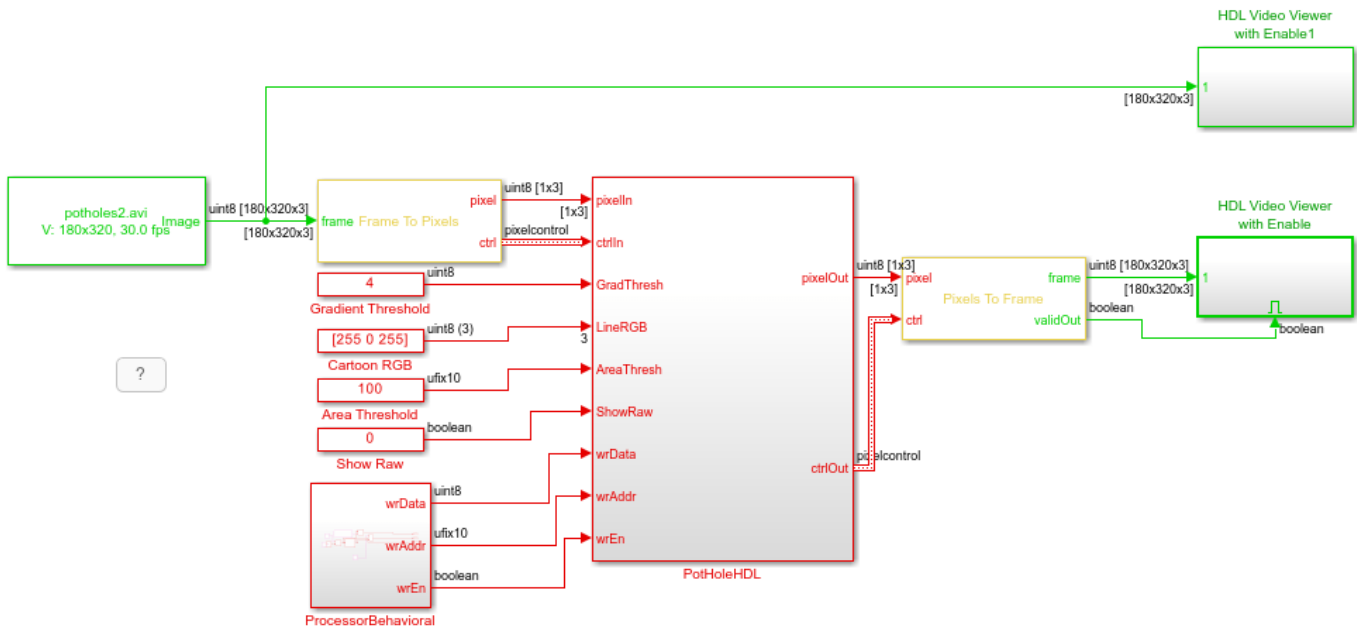
Copyright 2018-2022 The MathWorks, Inc.

Overview of the FPGA Subsystem

The PotHoleHDL subsystem converts the RGB input video to intensity, then performs bilateral filtering and edge detection. The TrapezoidalMask subsystem selects the roadway area. Then the design applies a morphological close and calculates centroid coordinates for all potential potholes. The detector selects the largest pothole in each frame and saves the center coordinates. The Pixel Stream Aligner matches the timing of the coordinates with the input stream. Finally, the Fiducial31x31 and the Overlay32x32 subsystems apply alpha channel overlays on the frame to add a pothole center marker and a text label.

```
open_system([modelname '/PotHoleHDL'],'force');
```
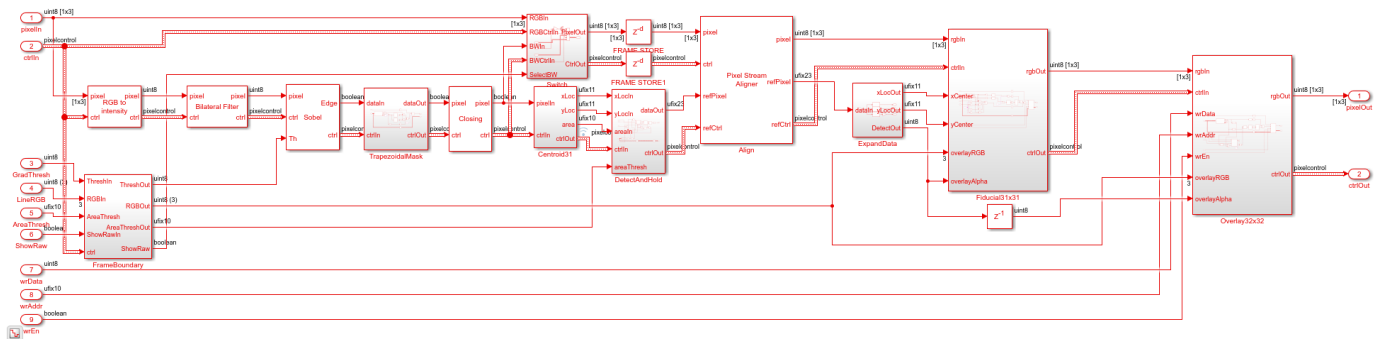

Input Parameter Values

The subsystem has four input parameters that can change while the system is running.

The gradient intensity parameter, Gradient Threshold, controls the edge detection part of the algorithm.

The Cartoon RGB parameter changes the color of the overlays, that is, the fiducial marker and the text.

The Area Threshold parameter sets the minimum number of marked pixels in the detection window in order for it to be classified as a pothole. If this value is too low, then linear cracks and other defects that are not road hazards will be detected. If it is too high then only the largest hazards will be detected.

The final parameter, Show Raw, allows you to debug the system more easily. It toggles the displayed image on which the overlays are drawn between the RGB input video and the binary image that the detector sees. Set this parameter to 1 to see how the detector is working.

All of these parameters work best if changes are only allowed on video frame boundaries. The FrameBoundary subsystem registers the parameters only on a valid start of frame.

```
open_system([modelname '/PotHoleHDL/FrameBoundary'],'force');
```
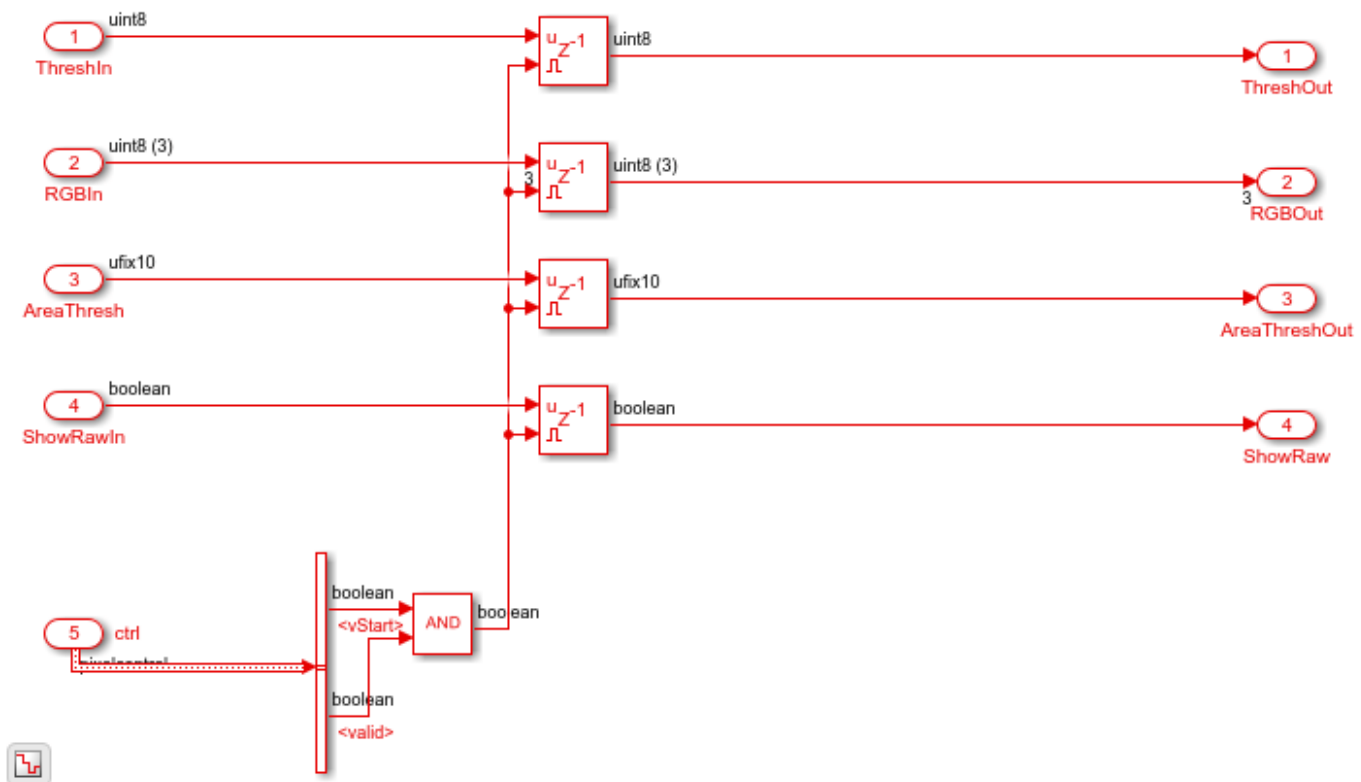

RGB to Intensity

The model splits the input RGB pixel stream so that a copy of the RGB stream continues toward the overlay blocks. The first step for the detector is to convert from RGB to intensity. Since the input data type for the RGB is `uint8`, the RGB to Intensity block automatically selects `uint8` as the output data type.

Bilateral Filter

The next step in the algorithm is to reduce high visual frequency noise and smaller road defects. There are many ways this can be accomplished but using a bilateral filter has the advantage of preserving edges while reducing the noise and smaller areas.

The Bilateral Filter block has parameters for the neighborhood size and two standard deviations, one for the spatial part of the filter and one for the intensity part of the filter. For this application a relatively large neighborhood of 9x9 works well. This model uses 3 and 0.75 for the standard deviations. You can experiment with these values later.

Sobel Edge Detection

The filtered image is then sent to the Sobel edge-detection block which finds the edges in the image and returns those edges that are stronger than the gradient threshold parameter. The output is a binary image. In your final application, this threshold can be set based on variables such as road conditions, weather, image brightness, etc. For this model, the threshold is an input parameter to the PotHoleHDL subsystem.

Trapezoidal Mask

From the binary edge image, you need to remove any edges that are not relevant to pothole detection. A good strategy is to use a mask that selects a polygonal region of interest and makes the area outside of that black. The model does not use a normal ROI block since that would remove the location context that you need later for the centroid calculation and labeling.

The order of operations also matters here because if you used the mask before edge detection, the edges of the mask would become strong lines that would result in false positives at the detector.

In the input video, the area in which the vehicle might encounter a pothole is limited to the roadway immediately in front of it and a trapezoidal section of roadway ahead. The exact coordinates depend on the camera mounting and lens. This example uses fixed coordinates for left-side top, right-side top, left-side bottom, and right-side bottom corners of the area. For this video, the top and bottom of the trapezoidal area are not parallel so this is not a true trapezoid.

The mask consists of straight lines between the corners, connecting left,right and top,bottom.

```
    ltc---rtc
   /         \
  /           \
 /             \
lbc------------rbc
```
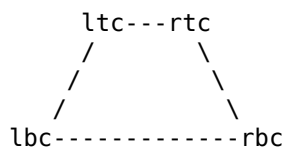
This example uses `polyfit` to determine a straight-line fit from corner to corner. For ease of implementation, the design calls `polyfit` with the vertical direction as the independent variable. This usage calculates $x = f(y)$ instead of the more usual $y = f(x)$. Using `polyfit` this way allows you to use a y-direction line counter as the input address of a lookup table of x-coordinates of the start (left) and end (right) of the area of interest on each line.

The lookup table is typically implemented in a BRAM in an FPGA, so it should be addressed with 0-based addressing. The model converts from MATLAB 1-based addressing to 0-based addressing just before the LUTs. To further reduce the size of the lookup table, the address is offset by the starting line of the trapezoid. In order to get good synthesis results, match typical block RAM registering in FPGAs by using a register after the lookup table. This register also adds some modest pipelining to the design.

For the 320x180 image:

```
raster = [320,180];
ltc = [155, 66];
lbc = [  1,140];
rtc = [155, 66];
rbc = [285,179];

% fit to x = f(y) for convenient LUT indexing
abl = polyfit([lbc(2),ltc(2)],[lbc(1),ltc(1)],1);  % left side
abr = polyfit([rbc(2),rtc(2)],[rbc(1),rtc(1)],1);  % right side
leftxstart = max(1,round((ltc(2):rbc(2))*abl(1)+abl(2)));
rightxend  = min(raster(1),round((ltc(2):rbc(2))*abr(1)+abr(2)));
startline  = min(ltc(2),rtc(2));
endline    = max(lbc(2),rbc(2));
% correct to zero-based addressing
leftxstart = leftxstart - 1;
rightxend = rightxend - 1;
startline = startline - 1;
endline = endline - 1;

open_system([modelname '/PotHoleHDL/TrapezoidalMask'],'force');
```
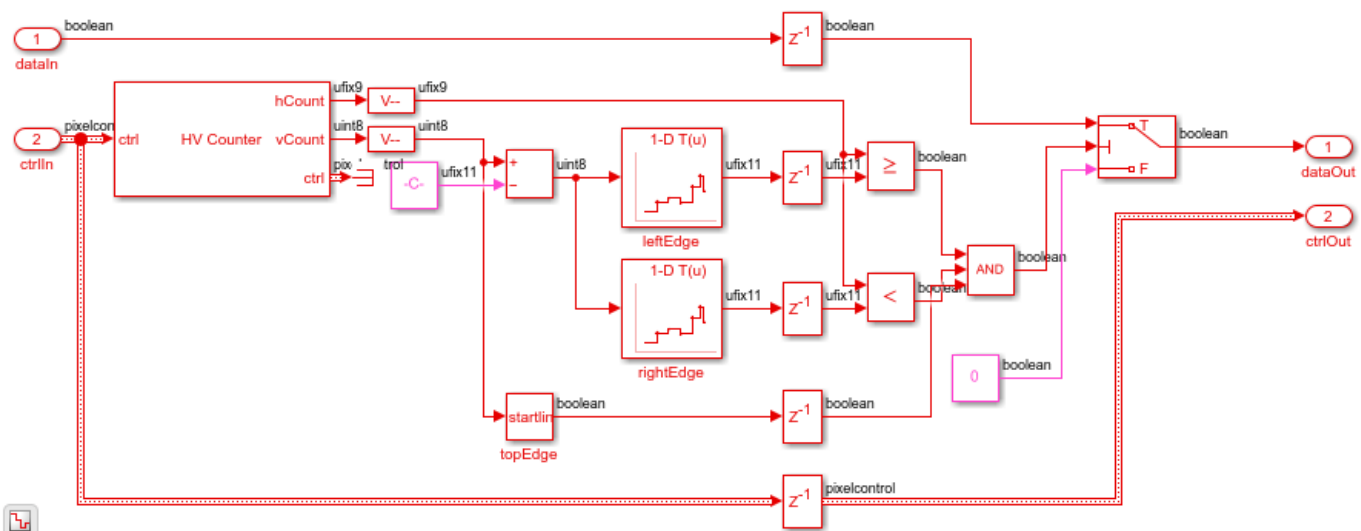

Morphological Closing

Next the design uses the Morphological Closing block to remove or close in small features. Closing works by first doing dilation and then erosion, and helps to remove small features that are not likely to be potholes. Specify a neighborhood on the block mask that determines how small or large a feature you want to remove. This model uses a 5x5 neighborhood, similar to a disk, so that small features are closed in.

Centroid

The centroid calculation finds the center of an active area. The design continuously computes the centroid of the marked area in each 31x31 pixel region. It only stores the center coordinates when the detected area is larger than an input parameter. This is a common difference between hardware

and software systems: when designing hardware for FPGAs it is often easier to compute continuously but only store the answer when you need it, as opposed to calling functions as-needed in software.

For a centroid calculation, you need to compute three things from the region of the image: the weighted sum of the pixels in the horizontal direction, the weighted sum in the vertical direction, and the overall sum of all the pixels which corresponds to the area of the marked portion of the region. The Line Buffer selects regions of 31x31 pixels, and returns them one column at a time. The algorithm uses the column to compute vertical weights, and total weights. For the horizontal weights, the design combines the columns to obtain a 31x31 kernel. You can choose the weights depending on what you want "center" to mean. This example uses -15:15 so that the center of the 31x31 region is (0,0) in the computed result.

The Vision HDL Toolbox blocks force the output data to zero when the output is not valid, as indicated in the pixelcontrol bus output. While not strictly required, this behavior makes testing and debugging much easier. To accomplish this behavior for the centroid results, the model uses Switch blocks with a Constant block set to 0.

Since you want the center of the detected region to be relative to the overall image coordinate system, add the horizontal and vertical pixel count to the calculated centroid.

```
open_system([modelname '/PotHoleHDL/Centroid31'],'force');
```
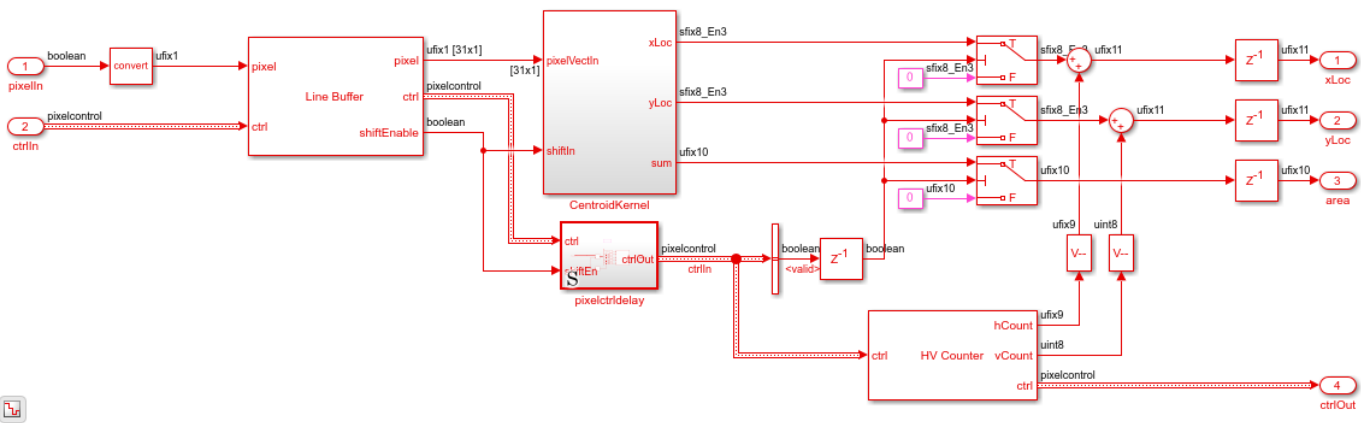


```
open_system([modelname '/PotHoleHDL/Centroid31/CentroidKernel'],'force');
```
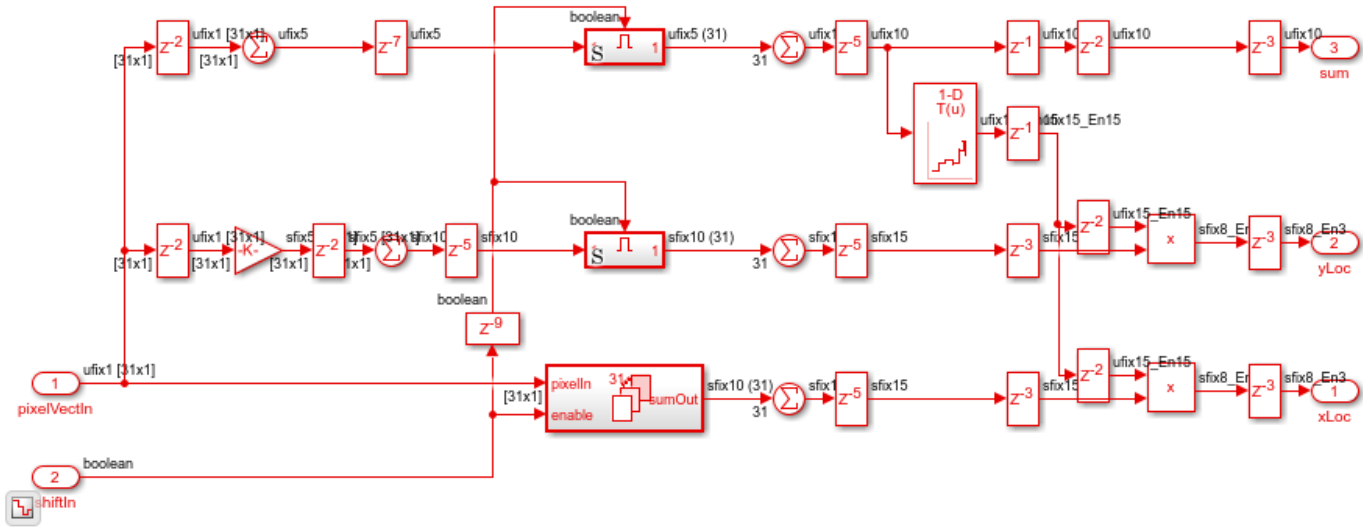
Detect and Hold

The detector operates on the total area sum from the centroid. The detector itself is very simple: compare the centroid area value to the threshold parameter, and find the largest area that is larger than the threshold. The model logic compares a stored area value to the current area value and stores a new area when the input is larger than the currently stored value. By using > or >= you can choose the earliest value over the threshold or the latest value over the threshold. The model stores the latest value because later values are closer to the camera and vehicle. When the detector stores a new winning area value, it also updates the X and Y centroid values that correspond to that area. These coordinates are then passed to the alignment and overlay parts of the subsystem.

To pass the X, Y, and valid indication to the alignment algorithm, pack the values into one 23-bit word. The model unpacks them once they are aligned in time with the input frames for overlay.

```
open_system([modelname '/PotHoleHDL/DetectAndHold'],'force');
```
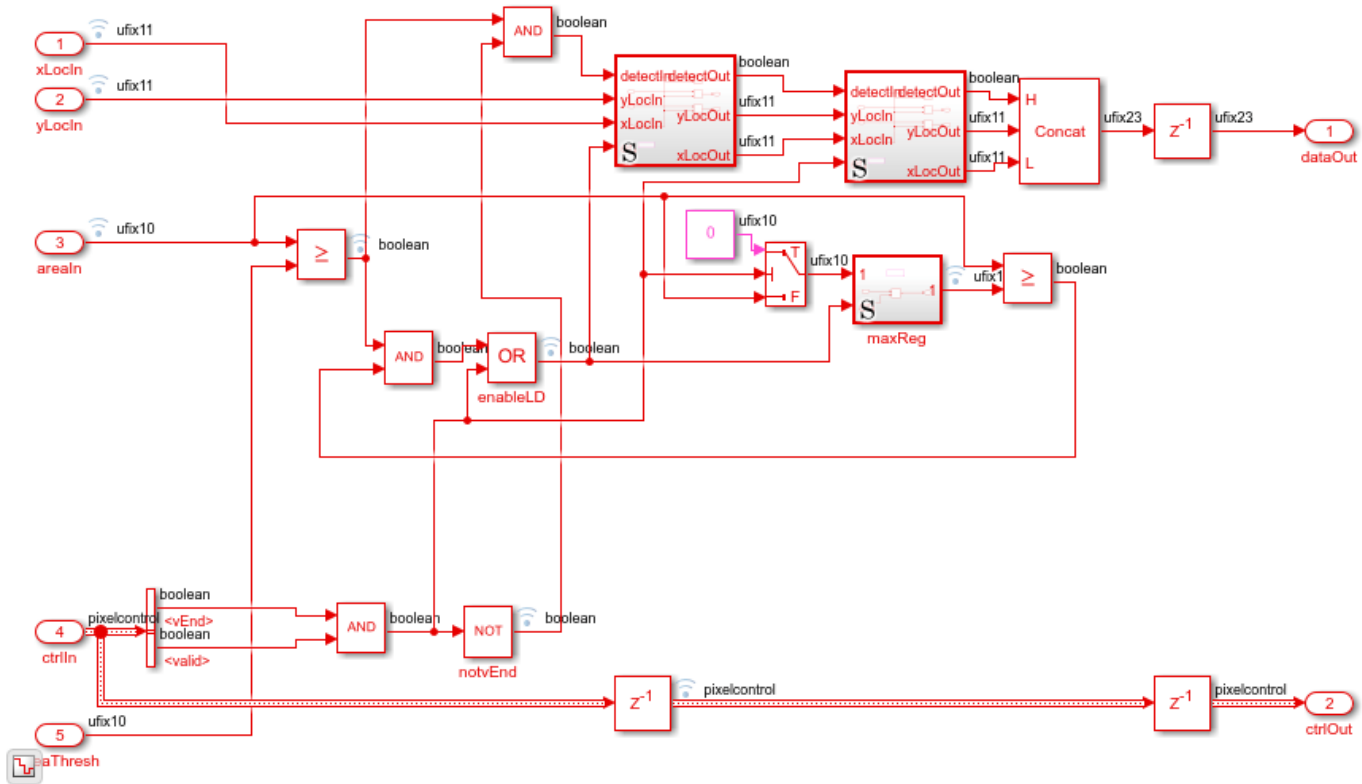
Pixel Stream Aligner

The Pixel Stream Aligner block takes the streaming information from the detector and sends it and the original RGB pixel stream to the overlay subsystems. The aligner compensates for the processing delay added by all the previous parts of the detection algorithm, without having to know anything about the latency of those blocks. If you later change a neighborhood size or add more processing, the aligner can compensate. If the total delay exceeds the **Maximum number of lines** parameter of the Pixel Stream Aligner block, adjust the parameter.

Fiducial Overlay

The fiducial marker is a square reticle represented as a 31-element array of 31-bit fixed-point numbers. This representation is convenient because a single read returns the whole word of overlay pixels for each line.

The diagram shows the overlay pattern by converting the fixed-point data to binary. This pattern can be anything you wish within the 31x31 size in this design.

```
load fiducialROM31x31.mat
crosshair = bin(fiducialROM);
crosshair(crosshair=='0') = ' ' %change '0' to space for better display


crosshair =

  31×31 char array

    '              1              '
```

```
'                 1                 '
'                 1                 '
'                 1                 '
'     111111111111111111111111     '
'     1           1           1     '
'     1           1           1     '
'     1           1           1     '
'     1           1           1     '
'     1           1           1     '
'     1           1           1     '
'     1           1           1     '
'     1                       1     '
'     1                       1     '
'     1                       1     '
'111111111111         111111111111'
'     1                       1     '
'     1                       1     '
'     1                       1     '
'     1           1           1     '
'     1           1           1     '
'     1           1           1     '
'     1           1           1     '
'     1           1           1     '
'     1           1           1     '
'     1           1           1     '
'     111111111111111111111111     '
'                 1                 '
'                 1                 '
'                 1                 '
'                 1                 '
```

The fiducial overlay subsystem has a horizontal and vertical counter with a set of four comparators that uses the center of the detected area as the center of the region for the marker. The marker data is used as a binary switch that turns on alpha channel overlay. The alpha value is a fixed transparency parameter applied as a gain on the binary Detect signal when it is unpacked, in the ExpandData subsystem.

```
open_system([modelname '/PotHoleHDL/Fiducial31x31'],'force');
```

Character Overlay

The character font ROM for the on-screen display stores data in a manner similar to the fiducial ROM described above. Each 16-bit fixed-point number represents 16 consecutive horizontal pixels. The character maps are 16x16.

Since the character data would typically be written by a CPU in ASCII, the simplest way is to store the character data under 8-bit ASCII addresses in a dual-port RAM. The font ROM stores ASCII characters 33 ("!") to 122 ("z"). The design offsets the address by 33.

The font ROM was constructed from a public domain fixed width font with a few edits to improve readability. As in the fiducial marker, the character ROM data is used as a binary switch that turns on alpha channel overlay. The character alpha value is a fixed transparency parameter applied as a gain on the Detect signal when it is unpacked, in the ExpandData subsystem.

To visualize the character B in the font ROM, display it in binary.

```
load charROM16x16.mat
letterB = bin(charROM16x16(529:544)); % character array
letterB(letterB=='0')=' ' % remove '0' chars for better display


letterB =

  16×16 char array

    '                '
    '   111111111    '
    '   11111111111  '
    '   111     111  '
    '   111      111 '
    '   111      111 '
```

```
'    111     111    '
'    1111111111     '
'    111111111      '
'    111     111    '
'    111       111  '
'    111       111  '
'    111       111  '
'    111      1111  '
'    11111111111    '
'    111111111      '
```

```
open_system([modelname '/PotHoleHDL/Overlay32x32'],'force');
```
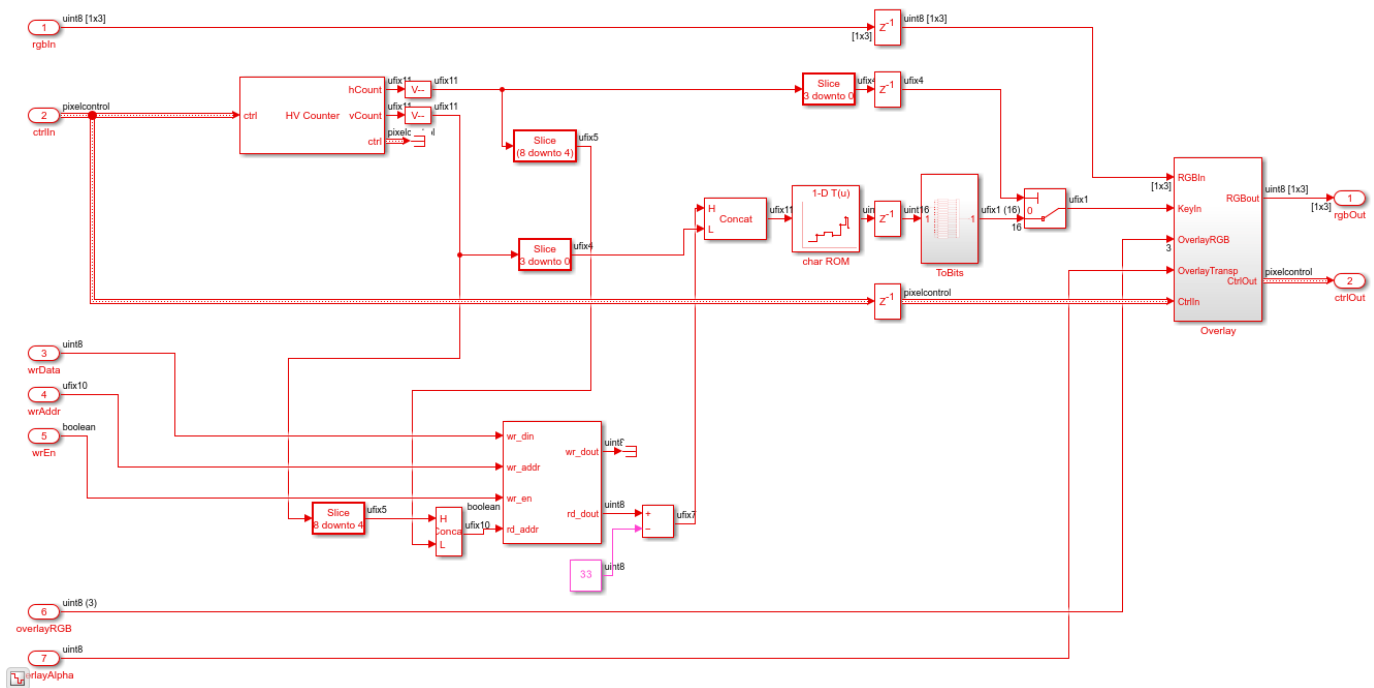

Viewing Detector Raw Image

When you work with a complicated algorithm, viewing intermediate steps in the processing can be very helpful for debugging and exploration. In this model, you can set the boolean `Show Raw` parameter to 1 (`true`) to display the result of morphological closing of the binary image, with the overlay of the detected results. To convert the binary image for use with the 8-bit RGB overlay, the model multiplies the binary value by 255 and uses that value on all three color channels.

HDL Code Generation

To check and generate the HDL code referenced in this example, you must have an HDL Coder™ license.

To generate the HDL code, use the following command.

```
makehdl('PotHoleHDLDetector/PotHoleHDL')
```

To generate the test bench, use the following command. Note that test bench generation takes a long time due to the large data size. You may want to reduce the simulation time before generating the test bench.

```
makehdltb('PotHoleHDLDetector/PotHoleHDL')
```

The part of this model that you can implement on an FPGA is the part between the Frame To Pixels and Pixels To Frame blocks. That is the subsystem called PotHoleHDL, which includes all the elements of the detector.

Simulation in an HDL Simulator

Now that you have HDL code, you can simulate it in your HDL Simulator. The automatically generated test bench allows you to prove that the Simulink simulation and the HDL simulation match.

Synthesis for an FPGA

You can also synthesize the generated HDL code in an FPGA synthesis tool, such as Xilinx Vivado. In a Virtex-7 FPGA (xc7v585tffg1157-1), the design achieves a clock rate of over 150 MHz.

The utilization report shows that the bilateral filter, pixel stream aligner, and centroid functions consume most of the resources in this design. The bilateral filter requires the most DSPs. The centroid implementation is quite efficient and uses only two DSPs. Centroid calculation also requires a reciprocal lookup table and so uses a large number of LUTs as memory.

Name ∧1	Slice LUTs (364200)	Slice Registers (728400)	F7 Muxes (182100)	Slice (91050)	LUT as Logic (364200)	LUT as Memory (111000)	LUT Flip Flop Pairs (364200)	Block RAM Tile (795)	DSPs (1260)
⊟ Ⓝ PotHoleHDL	18619	24540	92	7778	17012	1607	11170	305	92
⊞ Ⓐ u_Align (Align)	674	1065	90	519	646	28	149	96	0
⊞ Ⓐ u_Bilateral_Filter (Bilateral_...	6218	13879	0	3211	6070	148	5804	50	85
⊞ Ⓐ u_Centroid31 (Centroid31)	8836	6542	1	2794	7480	1356	3808	30.5	2
⊞ Ⓐ u_Closing (Closing)	1034	1057	0	400	998	36	619	8	0
⊞ Ⓐ u_Color_Space_Converte...	63	148	0	53	58	5	41	0	3
⊞ Ⓐ u_DetectAndHold (DetectA...	0	56	0	17	0	0	0	0	0
⊞ Ⓐ u_Edge_Detector (Edge_De...	569	926	0	282	544	25	376	2	2
⊞ Ⓐ u_Fiducial31x31 (Fiducial31...	266	163	0	147	262	4	82	0	0
Ⓐ u_FrameBoundary (FrameB...	77	86	0	85	77	0	2	0	0
⊞ Ⓐ u_Overlay32x32 (Overlay3...	495	178	1	193	490	5	103	1.5	0

Going Further

This example shows one possible implementation of an algorithm for detecting potholes. This design could be extended in the following ways :

- The gradient threshold could be computed from the average brightness using a gray-world model.
- The trapezoidal mask block could be made "steerable" by looking at the vehicle wheel position and adjusting the linear fit for the sloping sides of the mask.
- The detector could be made more robust by looking at the average brightness of the RGB or intensity image relative to the surrounding pavement since potholes are typically darker in intensity than the surrounding area.
- The visual frequency spectrogram of the pothole could also be used to look for specific types of surfaces in potholes.
- The detection area threshold value could be computed using average intensity in the trapezoidal roadway region.

- Multiple potholes could be detected in one frame by storing the top N responses rather than only the maximum detected response. The fiducial marker subsystem would need to be redesigned slightly to allow for overlapping markers.

Conclusion

This model shows how a pothole detection algorithm can be implemented in an FPGA. Many useful parts of this detector can be reused in other applications, such as the centroid block and the fiducial and character overlay blocks.

References

[1] Koch, Christian, and Ioannis Brilakis. "Pothole detection in asphalt pavement images." *Advanced Engineering Informatics* 25, no. 3 (2011): 507-15. doi:10.1016/j.aei.2011.01.002.

[2] Omanovic, Samir, Emir Buza, and Alvin Huseinovic. "Pothole Detection with Image Processing and Spectral Clustering." 2nd International Conference on Information Technology and Computer Networks (ICTN '13), Antalya, Turkey. October 2013.

Buffer Bursty Data Using Pixel Stream FIFO Block

This example shows how to interface with bursty pixel streams, such as those from DMA and Camera Link® sources, using the Pixel Stream FIFO block.

Overview

The DMACameraSourceHDL.slx system is shown below.
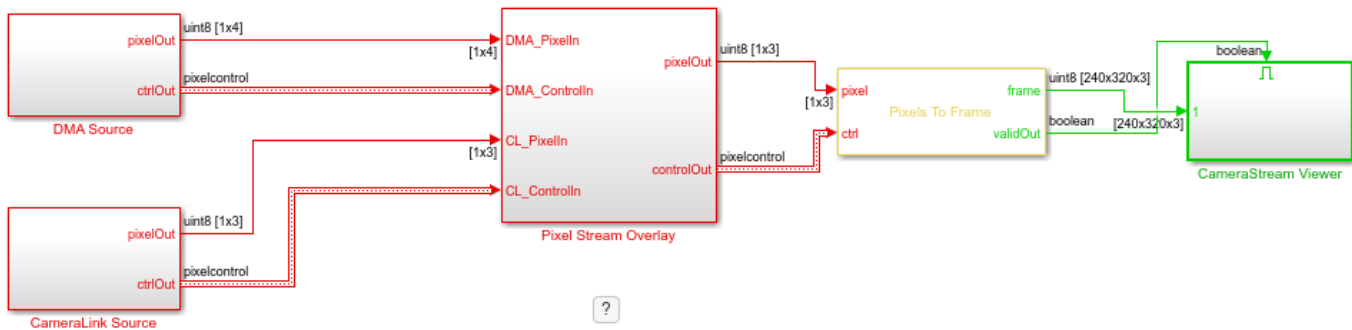
Copyright 2017 The MathWorks, Inc.

There are two pixel input streams - DMA source and Camera Link source. Input data for both sources is loaded from a `.mat` file, in the `InitFcn` callback. The DMA source models a non-contiguous stream of data arriving from off-chip memory. The pixels arrive in short bursts of random length, with random gaps between bursts. This can occur when there is contention on the DMA source and so it is not possible to stream pixels continuously from off-chip memory. The Camera Link source models the case when the camera is streaming an image of a lower resolution than the maximum permitted by the pixel clock and therefore will leave regular gaps between valid pixels. This spacing allows for streaming of multiple resolution images using a common clock, via strobing of validIn.

The Camera Link source models the incoming video stream from the sensor. The DMA source models a video stream from a frame buffer in which previous frame data has been processed in order to produce an alpha channel, allowing for blending of previous frame data with the current stream.

The **Pixel Stream Overlay** subsystem is shown in the diagram below. You can generate HDL code from this subsystem.
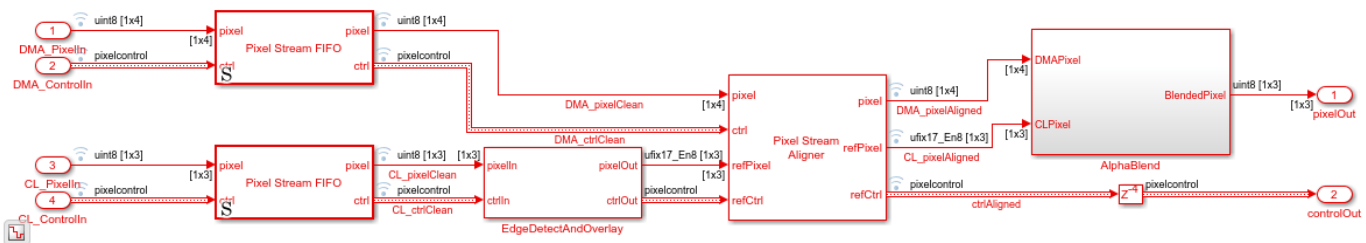
There are four main processing stages in the model - buffering of input data to remove burstiness, edge detection and overlay on Camera Link stream, alignment of pixel streams, and alpha blending of DMA stream onto Camera Link stream.

Pixel Stream Buffering

The Pixel Stream FIFO blocks buffer the input data as it is streamed into the model. The Pixel Stream FIFO is a masked subsystem. Looking into the Pixel Stream FIFO, we can see that it consists of a Memory Controller, Read and Write counters and two RAMs. One RAM stores the incoming pixel stream, and the other stores the incoming control signal stream. Once a full line has been buffered in RAM, the line is output continuously, removing any bursty behavior seen at input.
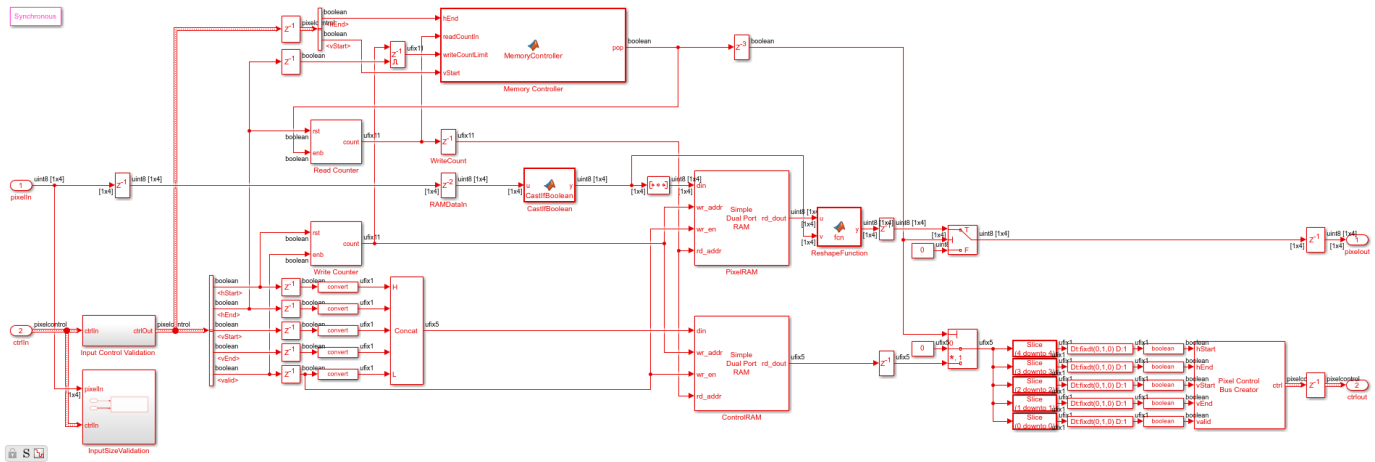
This waveform illustrates the difference in the pixel control signals after the Pixel Stream FIFO. The input `valid` signal, `DMA_ControlIn(5)`, shows short bursts of valid pixels, while the output `valid`, `DMA_ctrlClean(5)`, shows a continuous line of valid pixels. The total cycles in each line, shown by the time between `hStart` assertions, remains the same.
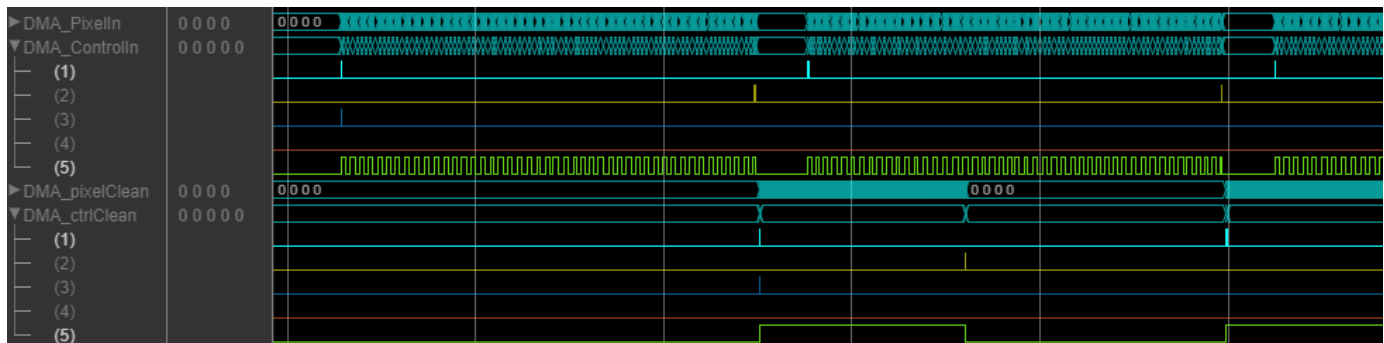
Edge Detection and Overlay on Camera Link Stream

To further differentiate the pixel streams, the Camera Link stream has an edge detection and overlay section. The pixel stream is first pre-processed by the Bilateral Filter block. This block smooths the image while preserving edges, and so it is a good choice for noise suppression prior to edge detection. The Edge Detector block detects edges using the Sobel method. The edges are then thinned using a [2x2] erosion operation. The thinned edge image is overlaid onto the original Camera Link image.
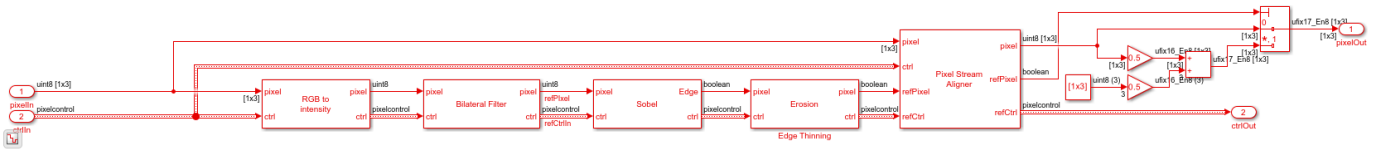
Pixel Stream Alignment

The Camera Link and DMA pixel streams must now be aligned to account for algorithmic delay in the data path. Aligning the pixel streams is straightforward using the Pixel Stream Aligner block.

Alpha Blending

The DMA input stream is a [1x4] vector whereas the Camera Link input is a [1x3] vector. The extra column in the DMA input is used to store the alpha channel information. The alpha channel represents the amount by which each of the pixels from the DMA source should be blended with the incoming Camera Link stream.
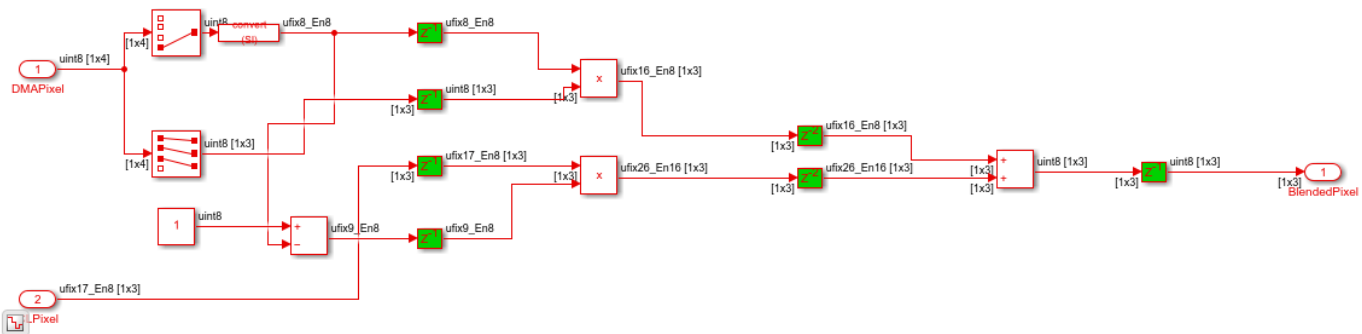
Results of the Simulation

The output video stream shows the DMA stream alpha blended onto the Camera Link input. The magenta colored overlay indicates the edges detected in the incoming Camera Link stream.
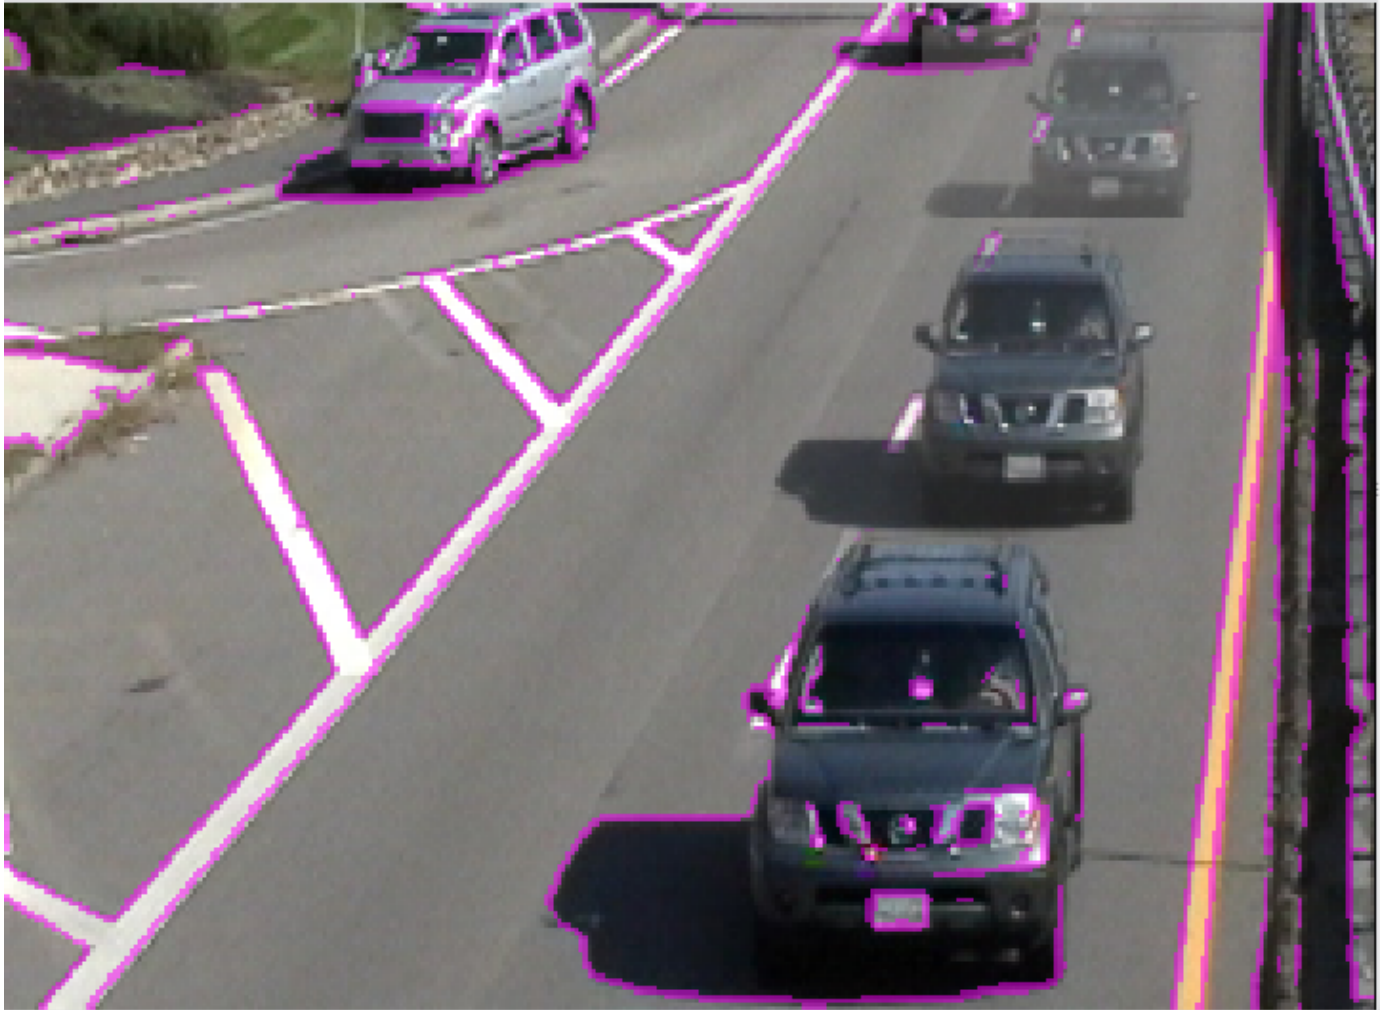
Generate HDL Code and Verify Its Behavior

To check and generate the HDL code referenced in this example, you must have an HDL Coder™ license.

To generate the HDL code, use the following command:

```
makehdl('DMACameraSourceHDL/Pixel Stream Overlay');
```

To generate an HDL test bench, use the following command:

```
makehdltb('DMACameraSourceHDL/Pixel Stream Overlay');
```

Using the Line Buffer to Create Efficient Separable Filters

This example shows how to design and implement a separable image filter, which uses fewer hardware resources than a traditional 2-D image filter.

Traditional image filters operate on a region of pixels and compute the resulting value of one central pixel using a two-dimensional filter kernel which is a matrix that represents the coefficients of the filter. Each coefficient is multiplied with its corresponding pixel and the result is summed to form the value. The region is then moved by one pixel and the next value is computed.

A separable filter is simple in concept: if the two-dimensional filter kernel can be factored into a horizontal component and a vertical component, then each direction can be computed separately using one-dimensional filters. This factorization can only be done for certain types of filter kernels. These kernels are called separable since the parts can be separated. Deciding which kernels are separable and which are not is easy using linear algebra in MATLAB. Mathematically, the two 1-D filters convolve to equal the original 2-D filter kernel, but a separable filter implementation often saves hardware resources.

Introduction

The SeparableFilterHDL.slx system is shown below. The SeptFiltHDL subsystem contains the separable filter, and also an Image Filter block implementation of the equivalent 2-D kernel as a reference.

```
modelname = 'SeparableFilterHDL';
open_system(modelname);
set_param(modelname,'SampleTimeColors','on');
set_param(modelname,'SimulationCommand','Update');
set_param(modelname,'Open','on');
set(allchild(0),'Visible','off');
```
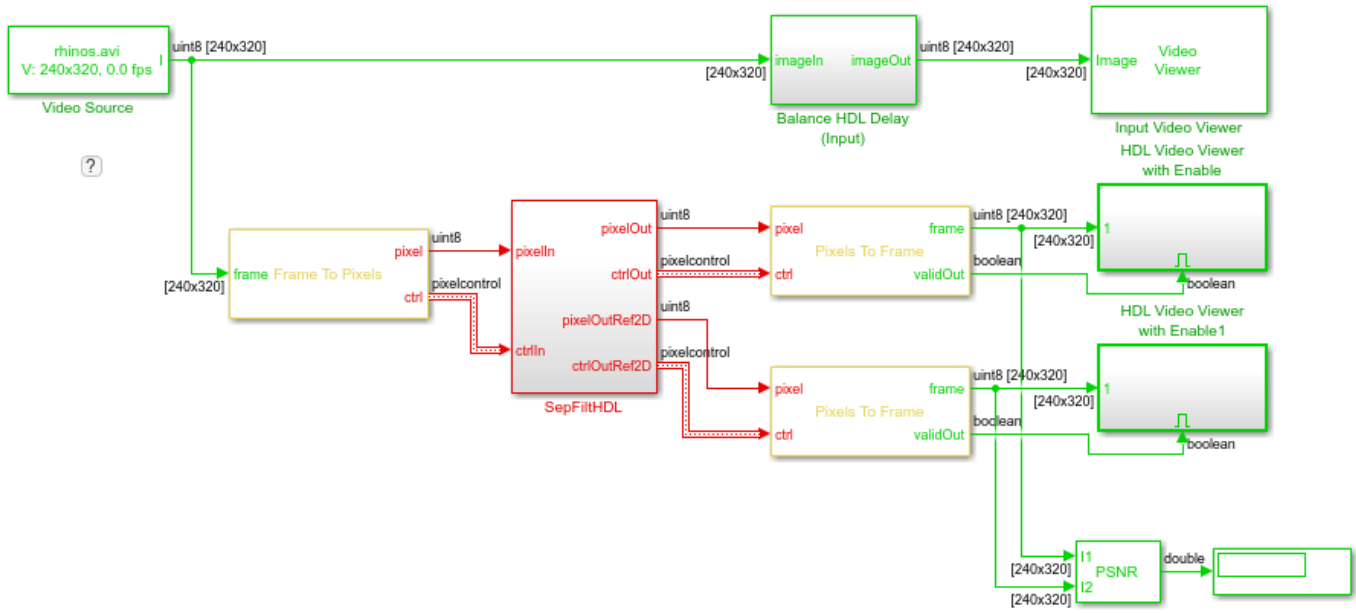
Copyright 2017 The MathWorks, Inc.

Determine Separable Filter Coefficients

Start by deciding what the purpose of your filter will be and compute the kernel. This example uses a Gaussian filter of size 5x5 with a standard deviation of 0.75. This filter has a blurring effect on images and is often used to remove noise and small details before other filtering operations, such as edge detection. Notice that the Gaussian filter kernel is circularly symmetric about the center.

```
Hg = fspecial('gaussian',[5,5],0.75)
```

```
Hg =

    0.0002    0.0033    0.0081    0.0033    0.0002
    0.0033    0.0479    0.1164    0.0479    0.0033
    0.0081    0.1164    0.2831    0.1164    0.0081
    0.0033    0.0479    0.1164    0.0479    0.0033
    0.0002    0.0033    0.0081    0.0033    0.0002
```

To check if the kernel is separable, compute its rank, which is an estimate of the number of linearly independent rows or columns in the kernel. If `rank` returns 1, then the rows and columns are related linearly and the kernel can be separated into its horizontal and vertical components.

```
rankHg = rank(Hg)
```

```
rankHg =

     1
```

To separate the kernel, use the `svd` function to perform singular value decomposition. The `svd` function returns three matrices, `[U,S,V]`, such that U*S*V' returns the original kernel, `Hg`. Since the kernel is rank 1, `S` contains only one non-zero element. The components of the separated filter are the first column of each of `U` and `V`, and the singular value split between the two vectors. To split the singular value, multiply both vectors with the square root of `S`. You must reshape `V` so that `Hh` is a horizontal, or row, vector.

For more information on filter separability, refer to the links at the bottom of this example.

```
[U,S,V]=svd(Hg)
Hv=abs(U(:,1)*sqrt(S(1,1)))
Hh=abs(V(:,1)'*sqrt(S(1,1)))
```

```
U =

   -0.0247   -0.9912   -0.1110    0.0673   -0.0000
   -0.3552    0.0737   -0.6054   -0.0431    0.7071
   -0.8640   -0.0301    0.4988    0.0619   -0.0000
   -0.3552    0.0737   -0.6054   -0.0431   -0.7071
   -0.0247   -0.0754    0.0761   -0.9939    0.0000


S =

    0.3793         0         0         0         0
         0    0.0000         0         0         0
         0         0    0.0000         0         0
         0         0         0    0.0000         0
         0         0         0         0    0.0000


V =

   -0.0247    0.1742    0.8521   -0.4929         0
   -0.3552    0.5980   -0.0716    0.1053    0.7071
   -0.8640   -0.4937    0.0199   -0.0968    0.0000
   -0.3552    0.5980   -0.0716    0.1053   -0.7071
   -0.0247   -0.1031    0.5131    0.8518   -0.0000


Hv =

    0.0152
    0.2188
    0.5321
    0.2188
    0.0152


Hh =

    0.0152    0.2188    0.5321    0.2188    0.0152
```

You can check your work by reconstructing the original kernel from these factored parts and see if they are the same to within floating-point precision. Compute the check kernel, `Hc`, and compare it to the original `Hg` using a tolerance.

```
Hc = Hv * Hh;
equalTest = all(Hc(:)-Hg(:) < 5*max(eps(Hg(:))))
```

```
equalTest =

  logical

   1
```

This result proves that Hv and Hh can be used to recreate the original filter kernel.

Fixed-Point Settings

For HDL code generation, you must set the filter coefficients to fixed-point data types. When picking fixed-point types, you must consider what happens to separability when you quantize the kernel.

First, quantize the entire kernel to a nominal data type. This example uses a 10-bit fixed-point number. Let the fixed-point tools select the best fraction length to represent the kernel values. Do the same conversion for the horizontal and vertical component vectors.

```
Hgfi = fi(Hg,0,10);
Hvfi = fi(Hv,0,10);
Hhfi = fi(Hh,0,10);
```

In this case, the best-precision 10-bit answer for Hg has 11 fractional bits, while Hv and Hh use only 10 fractional bits. This result makes sense since Hv and Hh are multiplied and added together to make the original kernel.

Now you must check if the quantized kernel is still rank 1. Since the rank and svd functions do not accept fixed-point types, you must convert back to doubles. This operation does not quantize the results, as long as the fixed-point type is smaller than 53 bits, which is the effective mantissa size of doubles.

```
rankDouble = rank(double(Hgfi))
```

```
rankDouble =

     3
```

This result shows that quantization can have a dramatic effect on the separability: since the rank is no longer 1, the quantized filter does not seem to be separable. For this particular filter kernel, you could experiment with the quantized word-length and discover that 51 bits of precision are needed in order for the rank function to return 1 after quantization. Actually, this result is overly conservative because of quantization of near-zero values within the rank function.

Instead of expanding the fixed-point type to 51 bits, add a tolerance argument to the rank function to limit the quantization effects.

```
rankDouble2048 = rank(double(Hgfi),1/2048)
```

```
rankDouble2048 =
```

```
    1
```

This result shows that the quantized kernel is rank 1 within an 11-bit fractional tolerance. So, the 11-bit separated coefficients are acceptable after all.

Another quantization concern is whether the fixed-point filter maintains flat field brightness. The filter coefficients must sum to 1.0 to maintain brightness levels. For a normalized Gaussian filter such as this example, the coefficient sum is always 1.0, but this sum can be moved away from 1.0 by fixed-point quantization. It can be critical in video and image processing to maintain exactly 1.0, depending on the application. If you imagine a chain of filters, each one of which raises the average level by around 1%, then the cumulative error can be large.

```
sumHg = sum( Hg(:) )
sumHgfi = sum( Hgfi(:) )


sumHg =

    1.0000


sumHgfi =

    1

          DataTypeMode: Fixed-point: binary point scaling
            Signedness: Unsigned
            WordLength: 15
         FractionLength: 11
```

In this case, the sums of the double-precision `Hg` and the fixed-point `Hgfi` are indeed 1.0. If maintaining brightness levels to absolute limits is important in your application, then you might have to manually adjust the filter coefficient values to maintain a sum of 1.

Finally, check that the combination of the quantized component filters still compares to the quantized kernel. By default, the `fi` function uses full precision on the arithmetic expression. Use convergent rounding since there are some coefficient values very near the rounding limit.

```
Hcfi = fi(Hvfi * Hhfi,0,10,'fimath',fimath('RoundingMethod','Convergent'));
equalTest = all( Hcfi(:)==Hgfi(:) )


equalTest =

  logical

    1
```

This result confirms that the fixed-point, separated coefficients achieve the same filter as the 2-D Gaussian kernel.

Implementing the Separable Filter

To see the separable filter implementation, open the Separable Filter subsystem that is inside the SepFiltHDL subsystem.

```
open_system([modelname '/SepFiltHDL/Separable Filter'],'force');
```
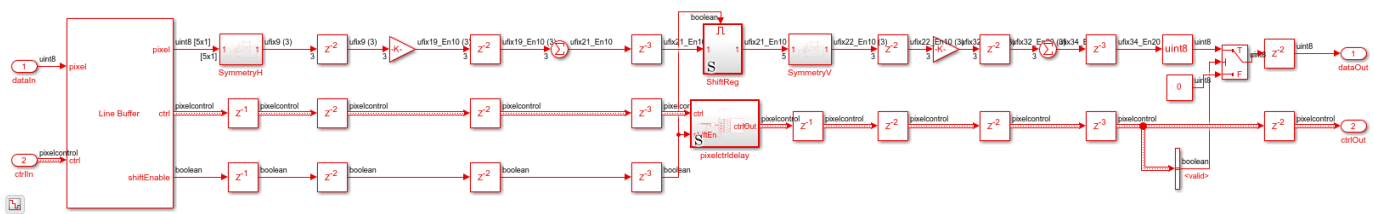

This subsystem selects vertical and horizontal vectors of pixels for filtering, and performs the filter operation.

The Line Buffer outputs a column of pixels for every time step of the filter. The Line Buffer also pads the edges of the image. This model uses **Padding method**: `Constant`, with a value of 0. The shiftEnable output signal is normally used to control a horizontal shift register to compile a 2-D pixel kernel. However, for a separable filter, you want to work in each direction separately. This model uses the output pixel column for the vertical filter, and uses the shiftEnable signal later to construct the horizontal pixel vector.

The separated horizontal and vertical filters are symmetric, so the model uses a pre-adder to reduce the number of multipliers even further. After the adder, a Gain block multiplies the column of pixels by the `Hv` vector. The Gain parameter is set to `Hv` and the parameter data type is `fixdt(0,10)`. The resulting output type in full-precision is `ufix18_En10`. Then a Sum block completes the vertical filter. The Sum block is configured in full-precision mode. The output is a scalar of `ufix21_En10` type.

There are many pipelining options you could choose, but since this design is simple, manual pipelining is quick and easy. Adding delays of 2 cycles before and after the Gain multipliers ensures good speed when synthesized for an FPGA. A delay of 3 cycles after the Sum allows for it to be sufficiently pipelined as well. The model balances these delays on the pixelcontrol bus and the shiftEnable signal before going to the horizontal dimension filter.

The best way to create a kernel-width shift register is to use a Tapped Delay block, which shifts in a scalar and outputs the register values as a vector. For good HDL synthesis results, use the Tapped Delay block inside an enabled subsystem, with the Synchronous marker block.

The output of the Tapped Delay subsystem is a vector of 5 horizontal pixels ready for the horizontal component of the separable filter. The model implements a similar symmetric pre-add and Gain block, this time with `Hh` as the parameter. Then, a Sum block and similar pipelining complete the horizontal filter. The final filtered pixel value is in the full-precision data type `ufix34_En20`.

Many times in image processing you would like to do full-precision or at least high-precision arithmetic operations, but then return to the original pixel input data type for the output. This subsystem returns to the original pixel type by using a Data Type Conversion block set to `uint8`, with `Nearest` rounding and saturation.

The Vision HDL Toolbox blocks force the output data to zero when the output is not valid, as indicated in the pixelcontrol bus output. While not strictly required, this behavior makes testing and debugging much easier. To accomplish this behavior, the model uses a Switch block with a Constant block set to 0.

Resource Comparison

The separable 5x5 filter implementation uses 3 multipliers in the vertical direction and 3 multipliers in the horizontal direction, for a total of 6 multipliers. A traditional image filter usually requires 25 multipliers for a 5x5 kernel. However, the Image Filter block takes advantage of any symmetry in the kernel. In this example the kernel has 8-way and 4-way symmetry, so the Image Filter only uses 5 multipliers. In general there are savings in multipliers when implementing a separable filter, but in this case the 2-D implementation is similar.

The separable filter uses 4 two-input adders in each direction, 2 for the pre-add plus 2 in the Sum, for a total of 8. The Image Filter requires 14 adders total, with 10 pre-add adders and 4 final adders. So there is a substantial saving in adders.

The Image Filter requires 25 registers for the shift register, while the separable filter uses only 5 registers for the shift register. Each adder also requires a pipeline register so that is 8 for the separable case and 14 for the traditional case. The number of multiplier pipeline registers scales depending on the number of multipliers.

The separable filter uses fewer adders and registers than the 2-D filter. The number of multipliers is similar between the two filters only because the 2-D implementation optimizes the symmetric coefficients.

Results of the Simulation

The resulting images from the simulation of the separable filter and the reference Image Filter are very similar. Using the fixed-point settings in this example, the difference between the separable filter and the reference filter never exceeds one bit. This difference is a 0.1% difference or greater than 54 dB PSNR between the filtered images overall.

HDL Code Generation

To check and generate the HDL code referenced in this example, you must have an HDL Coder™ license.

To generate the HDL code, use the following command.

```
makehdl('SeparableFilterHDL/SepFiltHDL')
```

To generate the test bench, use the following command. Note that test bench generation takes a long time due to the large data size. Reduce the simulation time before generating the test bench.

```
makehdltb('SeparableFilterHDL/SepFiltHDL')
```

The part of this model that you can implement on an FPGA is the part between the Frame To Pixels and Pixels To Frame blocks. The SepFiltHDL subsystem includes both the separable algorithm and the traditional 2-D implementation for comparison purposes.

Simulation in an HDL Simulator

Now that you have HDL code, you can simulate it in your HDL simulator. The automatically generated test bench allows you to prove that the Simulink simulation and the HDL simulation match.

Synthesis for an FPGA

You can also synthesize the generated HDL code in an FPGA synthesis tool, such as Xilinx Vivado. In a Virtex-7 FPGA (xc7v585tffg1157-1), the filter design achieves a clock rate of over 250 MHz.

The utilization report shows that the separable filter uses fewer resources than the traditional image filter. The difference in resource use is small due to the symmetry optimizations applied by the Image Filter block.

| Utilization - utilization_1 | | | | | | | | | | ? □ ↗ × |

Name	Slice LUTs (364200)	Slice Registers (728400)	Slice (91050)	LUT as Logic (364200)	LUT as Memory (111000)	LUT Flip Flop Pairs (364200)	Block RAM Tile (795)	DSPs (1260)	Bonded IOB (600)
SepFiltHDL	1660	2729	661	1508	152	1174	8	7	
u_Image_Filter (Image_Filter)	854	1500	370	793	61	623	4	4	
u_Separable_Filter (Separa...)	806	1198	333	715	91	536	4	3	

utilization_1

Going Further

The filter in this example is configured for Gaussian filtering but other types of filters are also separable, including some that are very useful. The mean filter, which has a kernel with coefficients that are all `1/N`, is always separable.

```
Hm = ones(3)./9
rank(Hm)


Hm =

    0.1111    0.1111    0.1111
    0.1111    0.1111    0.1111
    0.1111    0.1111    0.1111


ans =

    1

```

Or the Sobel edge-detection kernel:

```
Hs = [1 0 -1; 2 0 -2; 1 0 -1]
rank(Hs)


Hs =

    1    0    -1
    2    0    -2
    1    0    -1


ans =

    1

```

Or gradient kernels like this:

```
Hgrad = [1 2 3; 2 4 6; 3 6 9]
rank(Hgrad)
```

```
Hgrad =

    1     2     3
    2     4     6
    3     6     9


ans =

    1
```

Separability can also be applied to filters that do not use multiply-add, such as morphological filters where the operator is min or max.

Conclusion

You have used linear algebra to determine if a filter kernel is separable or not, and if it is, you learned how to separate the components using the `svd` function.

You explored the effects of fixed-point quantization and learned that it is important to work with precise values when calculating rank and singular values. You also learned about the importance of maintaining DC gain. Finally you learned why separable filters can be implemented more efficiently and how to calculate the savings.

References

[1] Eddins, S. "Separable convolution". Steve on Image Processing (October 4, 2006).

[2] Eddins, S. "Separable convolution: Part 2". Steve on Image Processing (November 28, 2006).

Image Pyramid

This example shows how to generate multi-level image pyramid pixel streams from an input stream. This model derives multiple pixel streams by downsampling the original image in both the horizontal and vertical directions, using Gaussian filtering. This type of filter avoids aliasing artifacts. The implementation uses an architecture suitable for FPGAs.

Image pyramid is used in many image processing applications such as image compression, object detection and recognition using techniques such as convolutional neural network (CNN) or aggregate channel features (ACF). Image pyramid is also similar to scale-space representation.

The example model takes a 240p video input and produces three output streams: 160x120, 80x60, and 40x30.

```
modelname = 'ImagePyramidHDL';
open_system(modelname);
set_param(modelname,'SampleTimeColors','on');
set_param(modelname,'SimulationCommand','Update');
set_param(modelname,'Open','on');
set(allchild(0),'Visible','off');
```
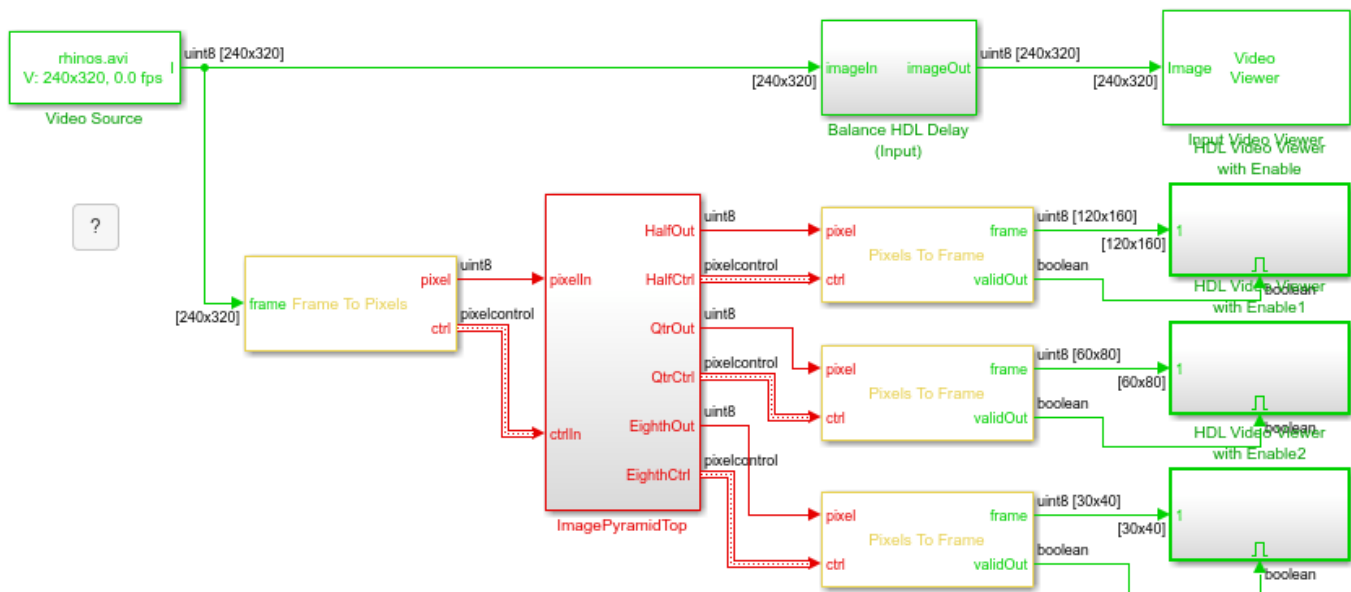

Copyright 2018 The MathWorks, Inc.

Each level of the pyramid contains a Line Buffer block and a downsampling filter.

```
open_system([modelname '/ImagePyramidTop/ResamplingPyramidFilter'],'force');
```
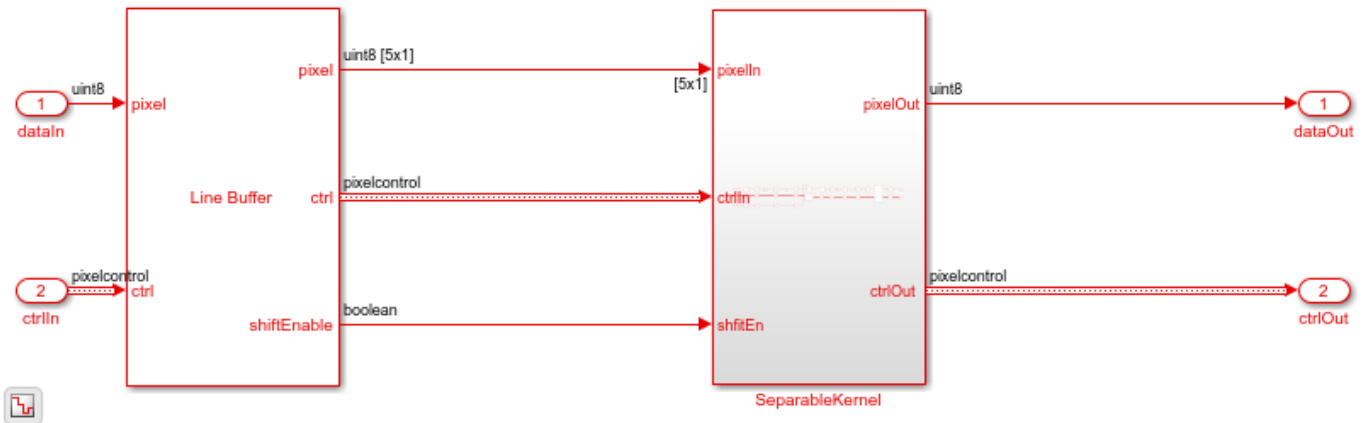
Filter Coefficients

The approximate Gaussian filter coefficients in [1] have been used in a number of image pyramid implementations. These coefficients are given by:

```
format long
Hh = [1 4 6 4 1]./16;
Hv = Hh';
Hg = Hv*Hh

Hg =

  Columns 1 through 3

    0.003906250000000   0.015625000000000   0.023437500000000
    0.015625000000000   0.062500000000000   0.093750000000000
    0.023437500000000   0.093750000000000   0.140625000000000
    0.015625000000000   0.062500000000000   0.093750000000000
    0.003906250000000   0.015625000000000   0.023437500000000

  Columns 4 through 5

    0.015625000000000   0.003906250000000
    0.062500000000000   0.015625000000000
    0.093750000000000   0.023437500000000
    0.062500000000000   0.015625000000000
    0.015625000000000   0.003906250000000
```

The results are similar to but not exactly the same as the Gaussian kernel with a 1.0817797 standard-deviation. So, Hg is an approximate Gaussian kernel.

```
Hf = fspecial('gaussian',5,1.0817797)

Hf =

  Columns 1 through 3

    0.004609023214619   0.016606534868404   0.025458671096979
    0.016606534868404   0.059834153028525   0.091728830511040
```

```
    0.025458671096979    0.091728830511040    0.140625009648116
    0.016606534868404    0.059834153028525    0.091728830511040
    0.004609023214619    0.016606534868404    0.025458671096979

  Columns 4 through 5

    0.016606534868404    0.004609023214619
    0.059834153028525    0.016606534868404
    0.091728830511040    0.025458671096979
    0.059834153028525    0.016606534868404
    0.016606534868404    0.004609023214619
```

The filter, `Hg`, is obviously separable, since it was constructed from horizontal and vertical vectors. Therefore, a separable filter implementation is a good choice. Many of the coefficient values are powers of two or a combination of only two powers of two. These values mean that the filter implementation can replace multiplication with shift and add techniques such as canonical signed digit (CSD). Each vector in the separable representation is also symmetric, so the filter implementation uses a symmetry pre-adder to further reduce the number of operations.

Downsampling

After low-pass filtering with the approximate Gaussian filter above, the model then downsamples the pixel stream by two in both the horizontal and vertical directions. This is primarily accomplished by alternating the valid signal every other pixel. The model also recreates the other `pixelcontrol` bus signals.

The model includes horizontal and vertical counters that compare the number of output pixels and lines with the mask parameters for active pixels and lines. The model uses these counts to recreate the end of line (hEnd) and end of frame (vEnd) signals.

After downsampling once, the `pixelcontrol` bus valid signal alternates high and then low every other pixel. After the second downsample, it alternates with a pattern of one valid pixel followed by three non-valid pixels. In some applications, you may want to collect all the valid pixels into a continuously valid period of time. The Pixel Stream FIFO block, used between downsample stages, produces continuous valid pixels for each line.

Each ResamplingPyramidFilter subsystem accepts parameters for the output frame size. These numbers must be integers, and a factor of two smaller than the input image. If the input number of pixels per line is odd rather than even, then round down to the next integer. For example, if the input size is 25 pixels per line, the requested output size must be 12 pixels per line.

Going Further

The Gaussian filter kernel used in traditional image pyramid is not the only low-pass filter that could be used. Using an edge-preserving low-pass filter, such as a bilateral filter, with different kernel sizes, would preserve more detail in the pyramid.

It is sometimes helpful to compute the difference between two levels of an image pyramid. This algorithm is called a Laplacian pyramid. The smaller level is upsampled to same size as the larger level, and filtered. The filter is usually a scaled version of the same approximate Gaussian filter used in this model. The difference between layers represents the information lost in the downsampling process. A Laplacian pyramid can be used for applications including coring for noise removal, compositing images taken at different times or with different focal lengths, and many others.

A potential limitation of this model is that there is fairly high latency between the output streams. This latency occurs because the second and third levels depend on the output from the previous level. This could be avoided by creating parallel filters operating on more lines. This example implements a 5-by-5 filter that stores 5 lines at each level. A lower latency parallel implementation requires 13 lines of storage for a two-level filter or 103 lines for a three-level filter. This is not generally a cost-effective trade-off.

On FPGAs, line buffer memories are typically implemented using block RAMs. Smaller memories can be implemented in the FPGA fabric, and are known as distributed RAMs. Your synthesis tool chooses block or distributed RAM depending on the resources of your device. As the line size becomes smaller due to downsampling, distributed RAMs can be more efficient. In this example, the Line Buffer blocks in each level reserve space for up to 2k pixels per line. This size is the default size for the Line Buffer, and accommodates up to 1080p format video. To target distributed RAMs, specify a small power of two for the **Line buffer size** parameter. In this example, you could set the line buffer sizes of the three levels to 256, 128, and 64.

References

[1] Burt, P., and E. Adelson. "The Laplacian Pyramid as a Compact Image Code."*IEEE Transactions on Communications* 31, no. 4 (April 1983): 532-40.

Stereo Disparity Using Semi-Global Block Matching

This example shows how to compute disparity between left and right stereo camera images using the Semi-Global Block Matching algorithm. This algorithm is suitable for implementation on an FPGA.

Distance estimation is an important measurement for applications in Automated Driving and Robotics. A cost-effective way of performing distance estimation is by using stereo camera vision. With a stereo camera, depth can be inferred from point correspondences using triangulation. Depth at any given point can be computed if the disparity at that point is known. Disparity measures the displacement of a point between two images. The higher the disparity, the closer the object.

This example computes disparity using the Semi-Global Block Matching (SGBM) method, similar to the `disparity` (Computer Vision Toolbox) function. The SGBM method is an intensity-based approach and generates a dense and smooth disparity map for good 3D reconstruction. However, it is highly compute-intensive and requires hardware acceleration using FPGAs or GPUs to obtain real-time performance.

The example model presented here is FPGA-hardware compatible, and can therefore provide real-time performance.

Introduction

Disparity estimation algorithms fall into two broad categories: local methods and global methods. Local methods evaluate one pixel at a time, considering only neighboring pixels. Global methods consider information that is available in the whole image. Local methods are poor at detecting sudden depth variation and occlusions, and hence global methods are preferred. Semi-global matching uses information from neighboring pixels in multiple directions to calculate the disparity of a pixel. Analysis in multiple directions results in a lot of computation. Instead of using the whole image, the disparity of a pixel can be calculated by considering a smaller block of pixels for ease of computation. Thus, the Semi-Global Block Matching (SGBM) algorithm uses block-based cost matching that is smoothed by path-wise information from multiple directions.

Using the block-based approach, this algorithm estimates approximate disparity of a pixel in the left image from the same pixel in the right image. More information about Stereo Vision is available here. Before going into the algorithm and implementation details, two important parameters need to be understood: Disparity Levels and Number of Directions.

Disparity Levels: Disparity levels is a parameter used to define the search space for matching. As shown in figure below, the algorithm searches for each pixel in the Left Image from among D pixels in the Right Image. The D values generated are D disparity levels for a pixel in Left Image. The first D columns of Left Image are unused because the corresponding pixels in Right Image are not available for comparison. In the figure, w represents the width of the image and h is the height of the image. For a given image resolution, increasing the disparity level reduces the minimum distance to detect depth. Increasing the disparity level also increases the computation load of the algorithm. At a given disparity level, increasing the image resolution increases the minimum distance to detect depth. Increasing the image resolution also increases the accuracy of depth estimation. The number of disparity levels are proportional to the input image resolution for detection of objects at the same depth. This example supports disparity levels from 8 to 128 (both values inclusive). **The explanation of the algorithm refers to 64 disparity levels.** The models provided in this example can accept input images of any resolution.
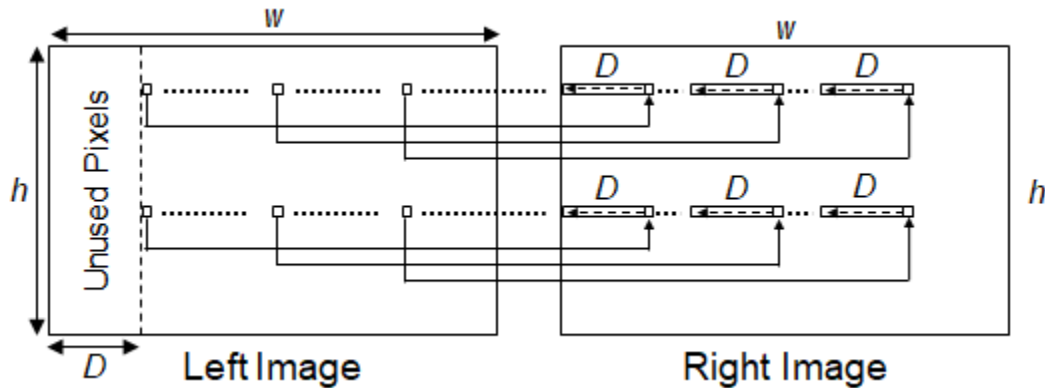
Number of Directions: In the SGBM algorithm, to optimize the cost function, the input image is considered from multiple directions. In general, accuracy of disparity result improves with increase in number of directions. This example analyzes five directions: left-to-right (A1), top-left-to-bottom-right (A2), top-to-bottom (A3), top-right-to-bottom-left (A4), and right-to-left (A5).

5 Directions
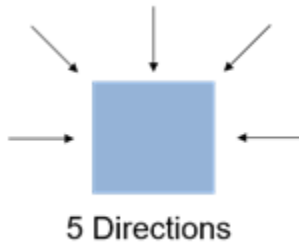
SGBM Algorithm

The SGBM algorithm takes a pair of rectified left and right images as input. The pixel data from the raw images may not have identical vertical coordinates because of slight variations in camera positions. Images need to be rectified before performing stereo matching to make all epi-polar lines parallel to the horizontal axis and match vertical coordinates of each corresponding pixel. For more details on rectification, please see `rectifyStereoImages` (Computer Vision Toolbox) function. The figure shows a block diagram of the SGBM algorithm, using five directions.
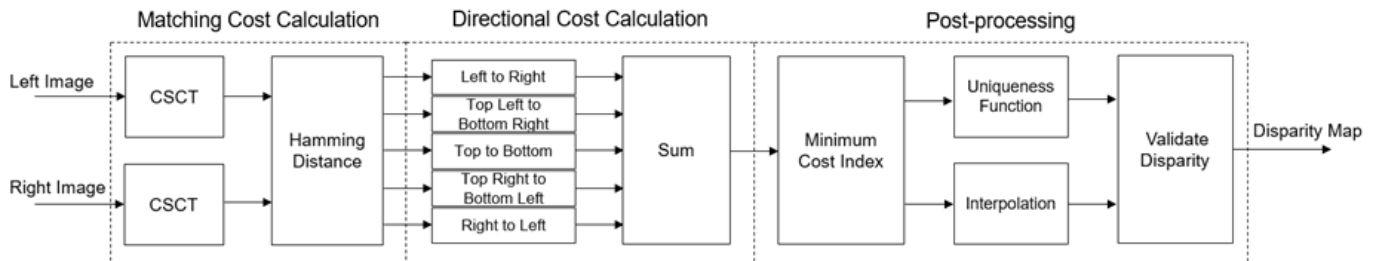
The SGBM algorithm implementation has three major modules: Matching Cost Calculation, Directional Cost Calculation and Post-processing.

Many methods have been explored in the literature for computing matching cost. This example implementation uses the census transform as explained in [2]. This module can be divided into two steps: Center-Symmetric Census Transform (CSCT) of left and right images and Hamming Distance computation. First, the model computes the CSCT on each of the left and right images using a sliding window. For a given pixel, a 9-by-7 pixel window is considered around it. CSCT for the center pixel in that window is estimated by comparing the value of each pixel with its corresponding center-symmetric counterpart in the window. If the pixel value is larger than its corresponding center-symmetric pixel, the result is 1, otherwise the result is 0. The figure shows an example 9-by-7 window. The center pixel number is 31. The 0th pixel is compared to the 62nd pixel (blue), the 1st pixel is compared to the 61st pixel (red), and so on, to generate 31 results. Each result a single bit output and the result of the whole window is arranged as a 31-bit number. This 31-bit number is the CSCT output for each pixel in both images.

0	1	2	3					
		31						
				59	60	61	62	

In the Hamming Distance module, the CSCT outputs of the left and right images are pixel-wise XOR'd and set bits are counted to generate the matching cost for each disparity level. To generate D disparity levels, D pixel-wise Hamming distance computation blocks are used. The matching cost for D disparity levels at a given pixel position, p, in the left image is computed by computing the Hamming distance with *(p to D+p)* pixel positions in the right image. The matching cost, *C(p,d)*, is computed at each pixel position, p, for each disparity level, d. The matching cost is not computed for pixel positions corresponding to the first D columns of the left image.

The second module of SGBM algorithm is directional cost estimation. In general, due to noise, the matching cost result is ambiguous and some wrong matches could have lower cost than correct ones. Therefore additional constraints are required to increase smoothness by penalizing changes of neighboring disparities. This constraint is realized by aggregating 1-D minimum cost paths from multiple directions. It is represented by aggregated cost from r directions at each pixel position, *S(p,d)*, as given by

$$S(p,d) = \sum_r L_r(p,d)$$

The 1-D minimum cost path for a given direction, *L_r(p,d)*, is computed as shown in the equation.

$$L_r(p,d) = C(p,d) + min(L_r(p-r,d), L_r(p-r,d-1) + P1, L_r(p-r,d+1) + P1, \min_i L_r(p-r,i) + P2) - \min_k L_r(p-r,k)$$

where

$L_r(p,d)$ = *current cost of pixel p and disparity d in direction r*

$C(p,d)$ = *matching cost at pixel p and disparity d*

$$L_r(p-r, d-1) \ = \ previous\ cost\ of\ pixel\ in\ r\ direction\ at\ disparity\ d-1$$

$$L_r(p-r, d+1) \ = \ previous\ cost\ of\ pixel\ in\ r\ direction\ at\ disparity\ d+1$$

$$\min_i L_r(p-r, i) \ = \ minimum\ cost\ of\ pixel\ in\ r\ direction\ for\ previous\ computation$$

$$P1,\ P2 \ = \ penalty\ for\ discontinuity$$

As mentioned earlier, this example uses five directions for disparity computation. Propagation in each direction is independent. The resulting disparities at each level from each direction are aggregated for each pixel. Total cost is the sum of the cost calculated for each direction.

The third module of SGBM algorithm is Post-processing. This module has three steps: minimum cost index calculation, interpolation, and a uniqueness function. Minimum cost index calculation finds the index corresponding to the minimum cost for a given pixel. Sub-pixel quadratic interpolation is applied on the index to resolve disparities at the sub-pixel level. The uniqueness function ensures reliability of the computed minimum disparity. A higher value of the uniqueness threshold marks more disparities unreliable. As a last step, the negative disparity values are invalidated and replaced with -1.

HDL Implementation

The figure below shows the overview of the example model. The blocks leftImage and rightImage import a stereo image pair as input to the algorithm. In the Input subsystem, the Frame To Pixels block converts input images from the leftImage and rightImage blocks to a pixel stream and accompanying control signals in a `pixelcontrol` bus. The pixel stream is passed as input to the SGBMHDLAlgorithm subsystem which contains three computation modules described above: matching cost calculation, directional cost calculation, and post-processing. The output of the SGBMHDLAlgorithm subsystem is a disparity value pixel stream. In the Output subsystem, the Pixels To Frame block converts the output to a matrix disparity map. The disparity map is displayed using the Video Viewer block.

```
modelname = 'SGBMDisparityExample';
open_system(modelname);
set_param(modelname,'SampleTimeColors','off');
set_param(modelname,'Open','on');
set_param(modelname,'SimulationCommand','Update');
set(allchild(0),'Visible','off');
```

FPGA Implementation of Stereo Disparity using SGBM
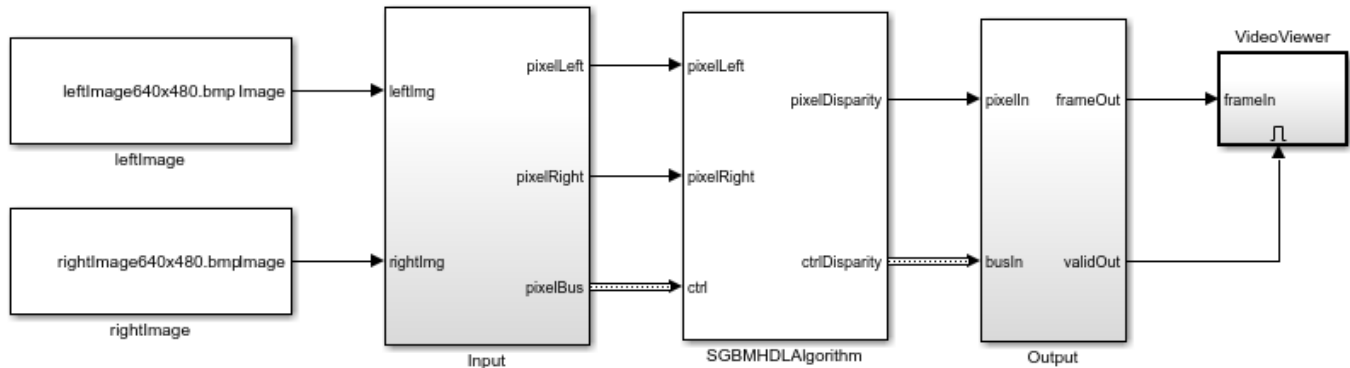
How to Use This Model?
Step 1 : Provide the path to rectified input images in the **leftImage** and **rightImage** blocks.
Step 2 : Run simulation and observe the output in the **VideoViewer**. The result is exported to the workspace with variable names `dispMap` and `dispMapValid`.

Parameter Selection:
Step 1 : Double-click the **SGBMHDLAlgorithm** block.
Step 2 : Enter Disparity Levels - an integer value between 8 and 128 [Default=64]
Step 3 : Enter a Uniqueness Threshold - an integer value between 5 and 15 [Default=5]

Copyright 2018-2022 The MathWorks, Inc.

Matching Cost Calculation

The matching cost calculation is again separated into two parts: CSCT computation and Hamming distance calculation. CSCT is calculated on each 9-by-7 pixel window by aligning each group of pixels for comparison using Tapped Delay (Simulink) blocks, For Each Subsystem (Simulink) blocks and buffers. The input pixels are padded with zeros to allow CSCT computation for the corner pixels. The resulting stream of pixels is passed to ctLogic subsystem. Figure below shows ctLogic subsystem which uses the Tapped Delay block to generate a group of pixels. The pixels are buffered for `imgColSize` cycles, where *imgColSize* is the number of pixels in an image line. A group of pixels that is aligned for comparison is generated from each row. The For Each block and Logical Operator block replicate the comparison logic for each pixel of the input vector size. To implement a 9-by-7 window, the model uses four such For Each blocks. The result generated by each For Each block is a vector which is further concatenated to form a vector of size 31-bits. After Bit Concat (HDL Coder) is used, the output data type is `uint5`. CSCT and zero-padding operations are performed separately on the left and right input images and the results are passed to the Hamming Distance subsystem.

```
open_system('SGBMDisparityExample/SGBMHDLAlgorithm/MatchingCost/CensusTransform/ctLogic','force')
```
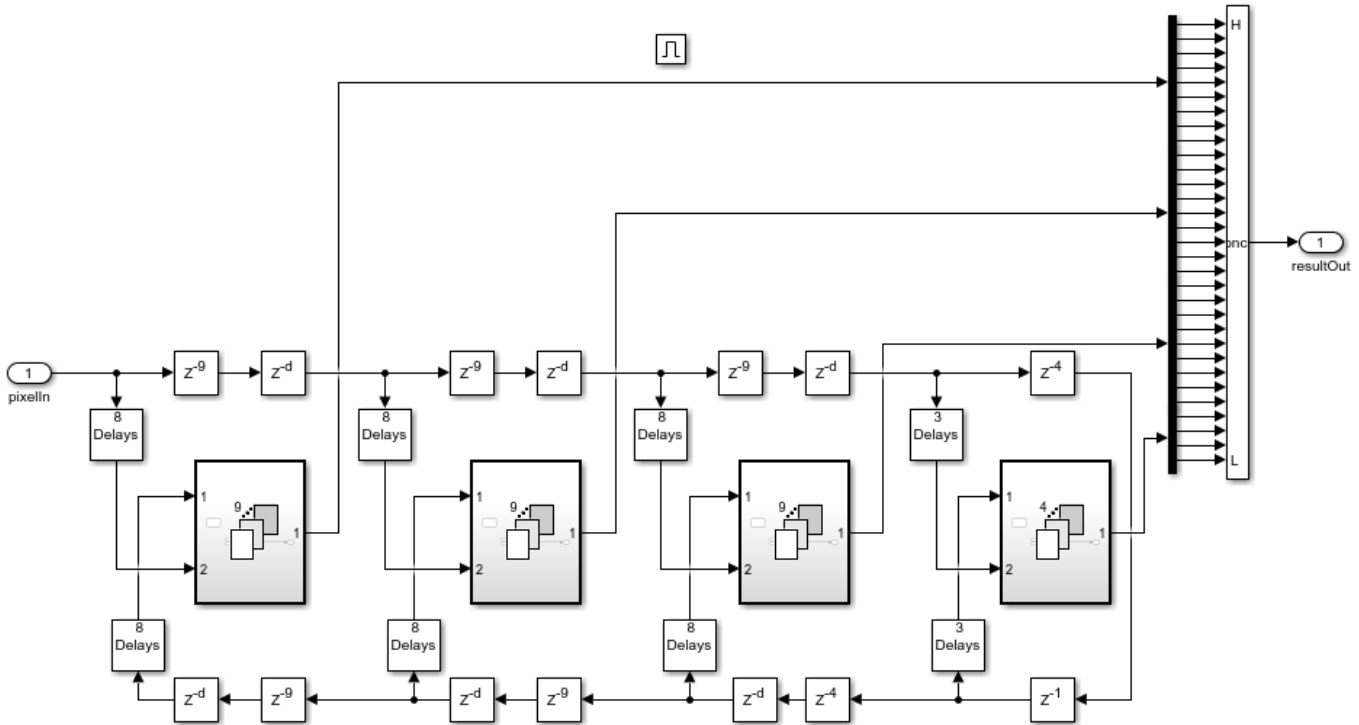
In the Hamming Distance subsystem, the 65th result of the left CSCT is XOR'd with the 65th to 2nd results of the right CSCT. The set bits are counted to obtain Hamming distance. This distance must be calculated for each disparity level. The right CSCT result is passed to the enabledTappedDelay subsystem to generate a group of pixels which is then XOR'd with the left CSCT result using For Each block. The For Each block also counts the set bits in the result. The For Each block replicates the Hamming distance calculation for each disparity level. The result is a vector, with 64 disparity levels corresponding to each pixel. This vector is the Matching Cost, and it is passed to the Directional Cost subsystem.

```
open_system('SGBMDisparityExample/SGBMHDLAlgorithm/MatchingCost/HammDistA','force');
```
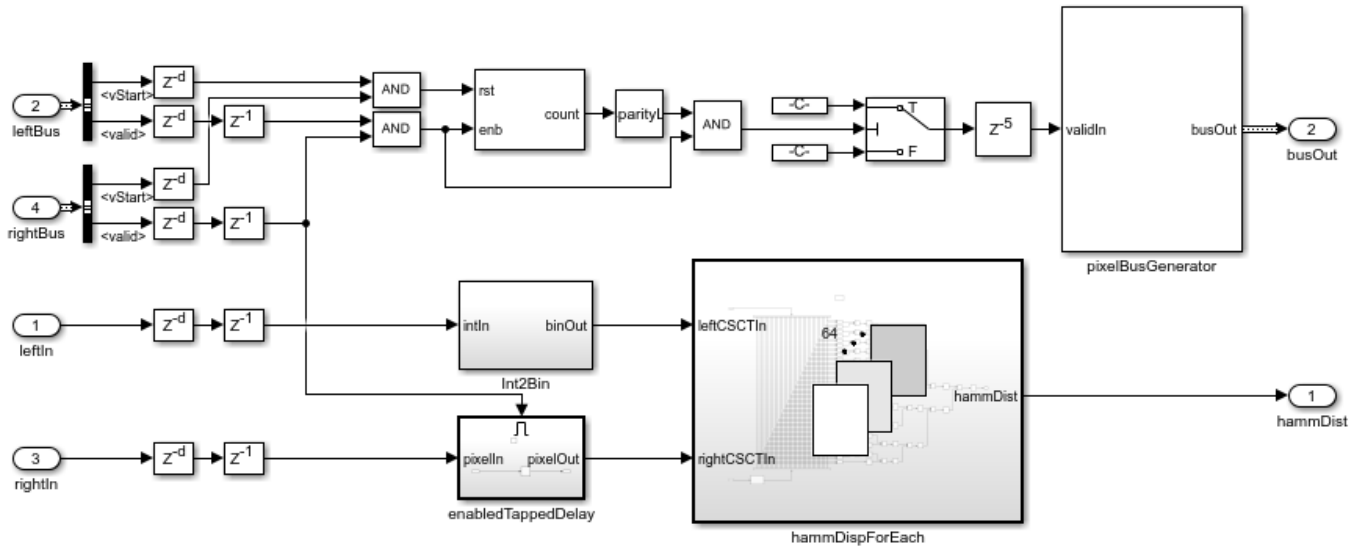
Directional Cost Calculation

The Directional Cost subsystem computes disparity at each pixel in multiple directions. The five directions used in the example are left-to-right (A1), top-left-to-bottom-right (A2), top-to-bottom (A3), top-right-to-bottom-left (A4), and right-to-left (A5). As the cost aggregation at each pixel in each direction is independent of each other, all five directions are implemented concurrently.

Each directional analysis is investigating the previous cost value with respect to the current cost value. The value of previous cost required to compute the current cost for each pixel depends on the direction under consideration. The figure below shows the position of the previous cost with respect to the current cost under computation, for all five directions.
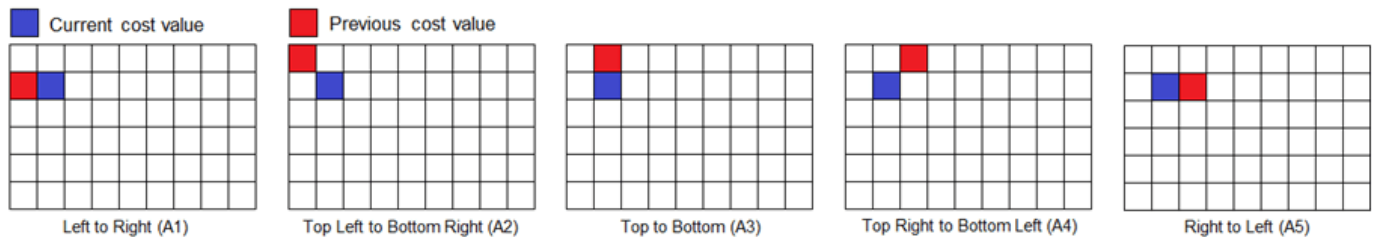
In the figure above, the blue box indicates the position of the current pixel for which current cost values are computed. The red box indicates the position of the previous cost values to be used for current cost computation. For A1, the current cost becomes the previous cost value for the next computation when traversing from left to right. Thus, the current cost value should be immediately fed back to compute the next current cost, as described in [3]. For A2, when traversing from left to right, current cost value should be used as the previous cost value after *imgColSize+1* cycles. Current cost values are hence buffered for cycles equal to *imgColSize+1* and then fed back to compute the next current cost.

Similarly, for A3 and A4, the current cost values are buffered for cycles equal to *imgColSize* and *imgColSize-1*, respectively. However, for A5, when traversing from left to right, the previous cost value is not available. Thus, the direction of traversal to compute A5 is reversed. This adjustment is

done by reversing the input pixels of each row. The current cost value then becomes the previous cost value for the next current cost computation, similar to A1.

The 1-D minimum cost path computes the current cost at *d* disparity position, using the Matching Cost value, the previous cost values at disparities *d-1*, *d*, and *d+1*, and the minimum of the previous cost values. The figure below shows the minimum cost path subsystem, which computes the current cost at a disparity position for a pixel.

`open_system('SGBMDisparityExample/SGBMHDLAlgorithm/DirectionalCost/LeftToRight/lrSubsystem/minCos`
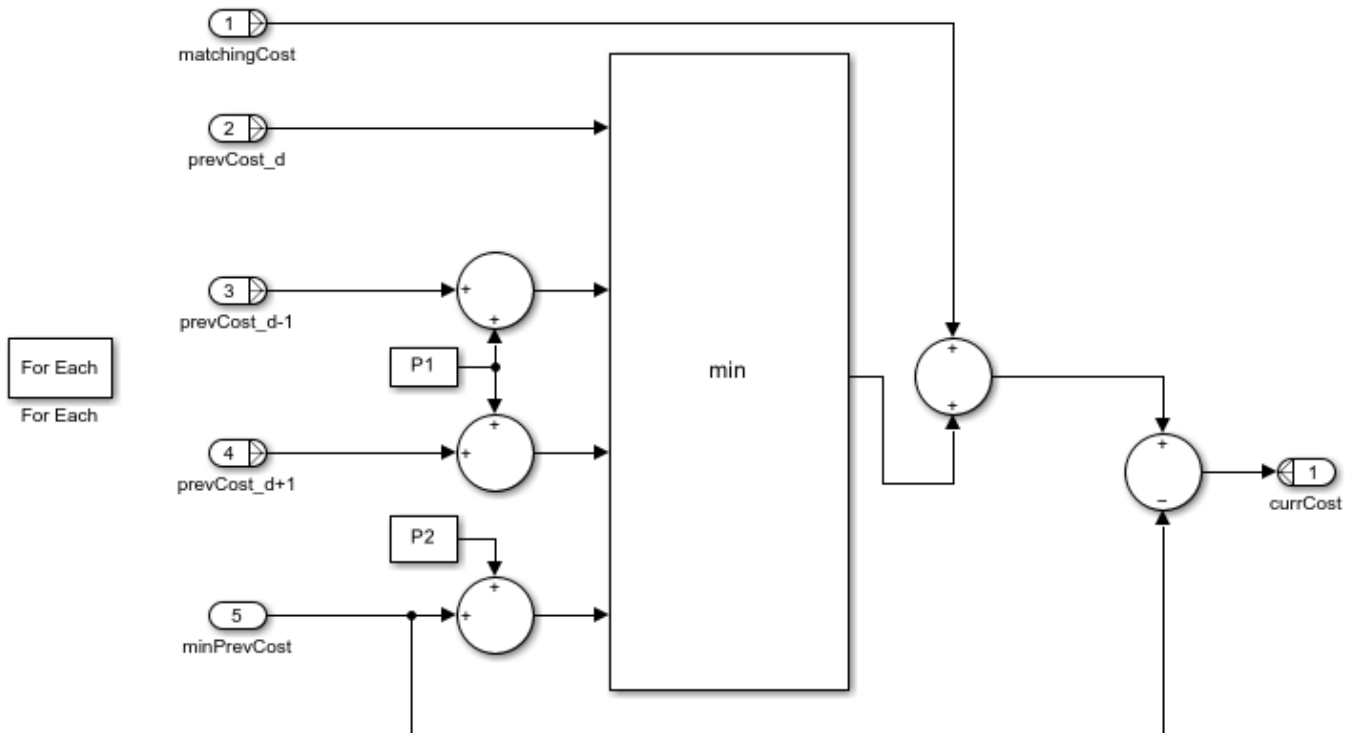
The For Each block is used to replicate the minimum cost path calculation for each disparity level, for each direction. The figure below shows the implementation of A1 for 64 disparity levels. As shown in the figure, 64 minimum cost path calculations are generated as represented by minCostPath subsystem. The matching cost is an input from the Hamming Distance subsystem. The current cost computed by the minCostPath subsystem is immediately fed back to itself as the previous cost values, for the next current cost computation. Thus, values for *prevCost_d* are now available. Values for *prevCost_d-1* are obtained by shifting the 1st to 63rd fed-back values to the 2nd to 64th positions. The d-1 subsystem contains a Selector (Simulink) block that shifts the position of the values, and fills in zero at the 1st position.

Similarly, values for *prevCost_d+1* position are obtained by shifting the 2nd to 64th feedback values to the 1st to 63rd position and inserting a zero at the 64th position. The current cost computed is also passed to the min block to compute the minimum value from the current cost values. This value is fed back to the *minPrevCost* input of the minCostPath subsystem. The next current cost is then computed by using the current cost values, acting as previous cost values, in the next cycle for A1. Since the minimum cost of disparity levels from the previous set is immediately needed for the current set, this feedback path is the critical path in the design.

`open_system('SGBMDisparityExample/SGBMHDLAlgorithm/DirectionalCost/LeftToRight/lrSubsystem','for`
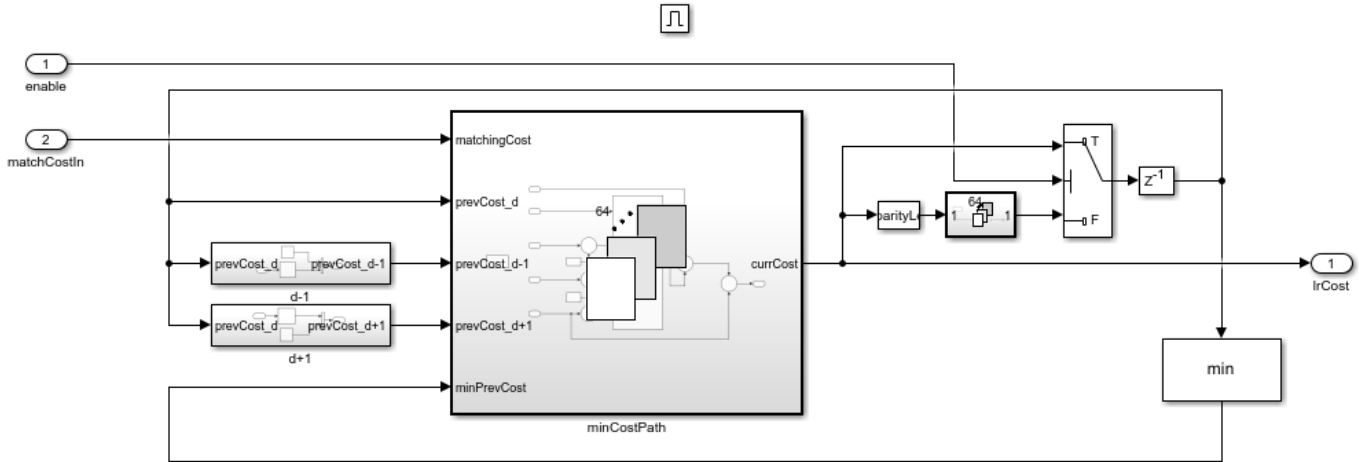
The current cost computations for A2, A3, and A4 are implemented in the same manner. Since the current cost value is not immediately required for these directions, there is a buffer in both feedback paths. This buffer prevents this feedback path from being the critical path in the design. The figure below shows the A3 implementation with a buffer in the feedback paths.

```
open_system('SGBMDisparityExample/SGBMHDLAlgorithm/DirectionalCost/TopToBottom/tbSubsystem','for
```
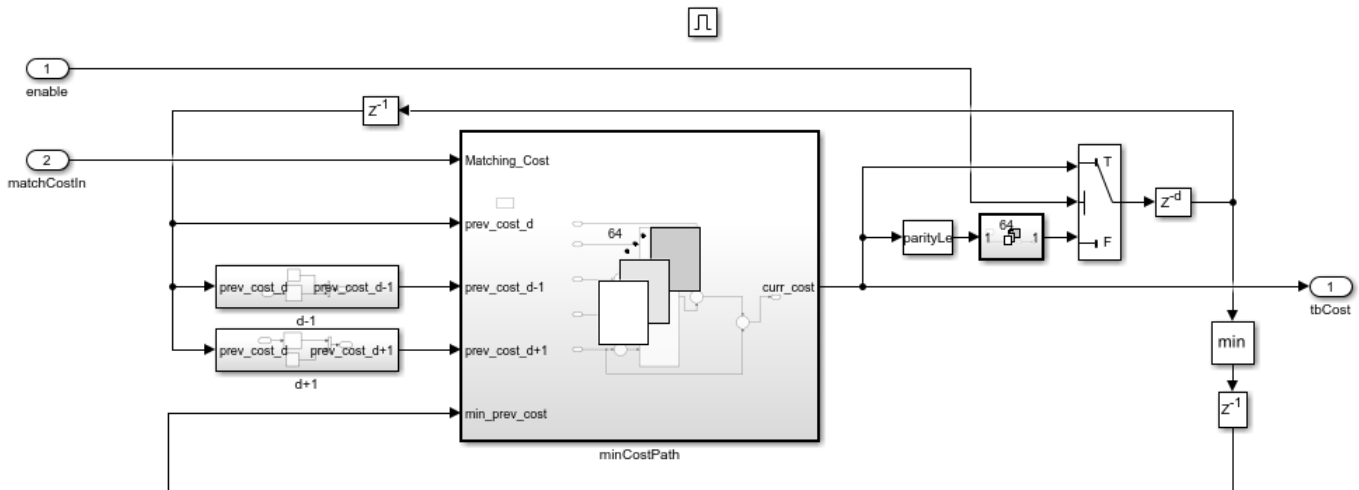

The current cost calculation for A5 has additional logic to reverse the rows at input and again reverse the rows at output to match the pixel positions for the total cost calculation. A single buffer of *imgColSize* cycles achieves this reversal. Since all directions are calculated concurrently, the time required to reverse the rows must be compensated for on the other paths. Delay equivalent to *2*imgColSize* cycles is introduced in the other four directions. To optimize resources, instead of buffering 64 values of matching cost for each pixel, the 31-bit result of CSCT is buffered. A separate Hamming Distance module is then required to compute matching cost for A5. This design reduces on-chip memory usage. The rows are reversed after the CSCT computation and matching cost is calculated using a separate Hamming Distance module that provides the Matching Cost input to A5. Also, dataAligner subsystem is used to remove data discontinuity for each row before passing it to Hamming Distance subsystems. This helps easy synchronization of data at time of aggregation. The

current cost obtained from all five directions at each pixel are aggregated to obtain the total cost at each pixel. The total cost is passed to the Post-processing subsystem.

Post-Processing

In the post-processing subsystem, the index of the minimum cost is calculated at each pixel position from 64 disparity levels by using Min blocks in a tree architecture. The index value obtained is the disparity of each pixel. Along with minimum cost index computation, the minimum cost value at the computed index, and the cost values at *index-1* and *index+1* are also computed. The Minimum_Cost_Index subsystem implements tree architecture to compute a minimum value from 128 values. 64 disparity values are padded with 64 more values to make a vector of 128 values. Minimum value is computed from this vector with 128 values. In case, a vector with 128 values is available no value is padded to a vector or in other words, vector is passed directly for minimum value calculation. Variant Subsystem, Variant Model, Variant Assembly Subsystem (Simulink) is used to select between logic using variant subsystem variables. Sub-pixel quadratic interpolation is then applied to the index to resolve disparity at sub-pixel level. Also, a uniqueness function is applied to the index calculated by min blocks, to ensure reliable disparity results. As a last step, invalid disparities are identified and replaced with -1.

Model Parameters

The model presented here takes disparity levels and uniqueness threshold as input parameters as shown in figure. Disparity levels is an integer value from 8 to 128 with the default value of 64. Higher value of disparity level reduces the minimum distance detected. Also, for larger input image size larger disparity level helps better detection of depth of object. The uniqueness threshold must be a positive integer value, between 0 and 100 with a typical range from 5 to 15. Lower value of uniqueness threshold marks more disparities reliable. The default value of uniqueness threshold is 5.
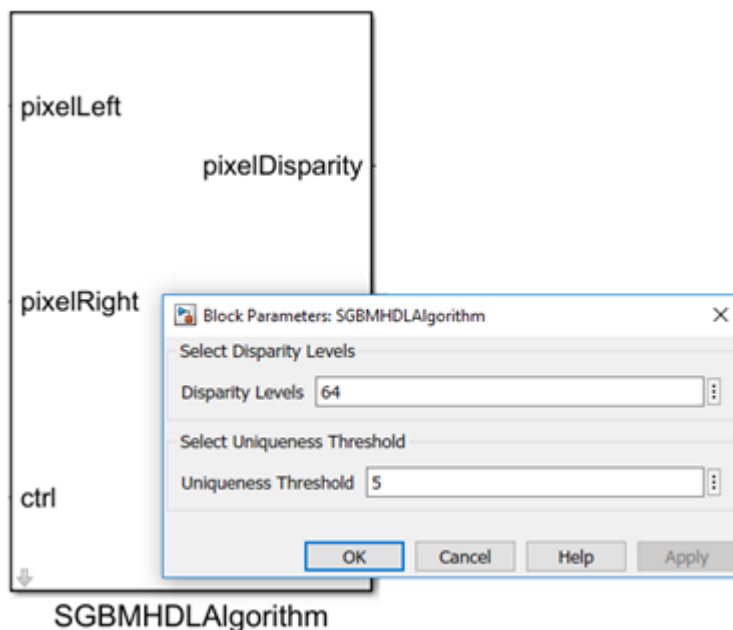
Simulation and Results

The example model can be simulated by specifying a path for the input images in the leftImage and rightImage blocks. The example uses sample images of size 640-by-480 pixels. The figure shows a

sample input image and the calculated disparity map. The model exports these calculated disparities and a corresponding valid signal to the MATLAB workspace, using variable names `dispMap` and `dispMapValid` respectively. The output disparity map is 576-by-480 pixels, since the first 64 columns are unused in the disparity computation. The unused pixels are padded with 0 in Output subsystem to generate output image of size 640-by-480 as shown in Video Viewer. A disparity map with colorbar is generated using the commands shown below. Higher disparity values in the result indicate that the object is nearer to the camera and lower disparity values indicate farther objects.

```
dispMapValid = find(dispMapValid == 1);
disparityMap = (reshape(dispMap(dispMapValid(1:imgRowSize*imgColSize),:),imgColSize,imgRowSize))
figure(); imagesc(disparityMap);
title('Disparity Map');
colormap jet; colorbar;
```


leftImage640x480

rightImage640x480

Video Viewer
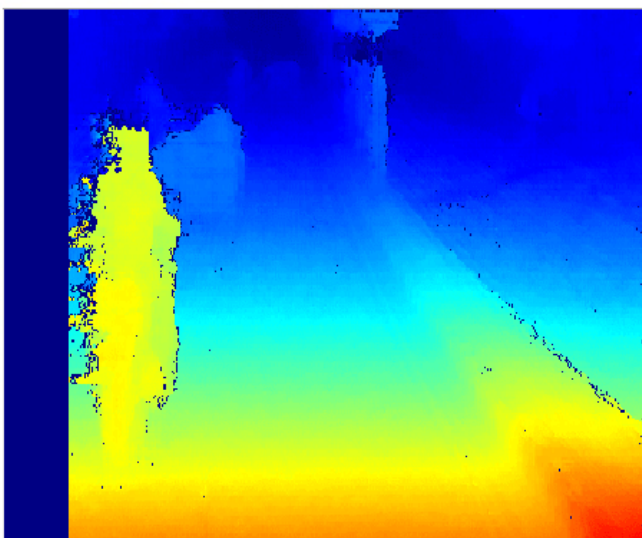
Disparity Map
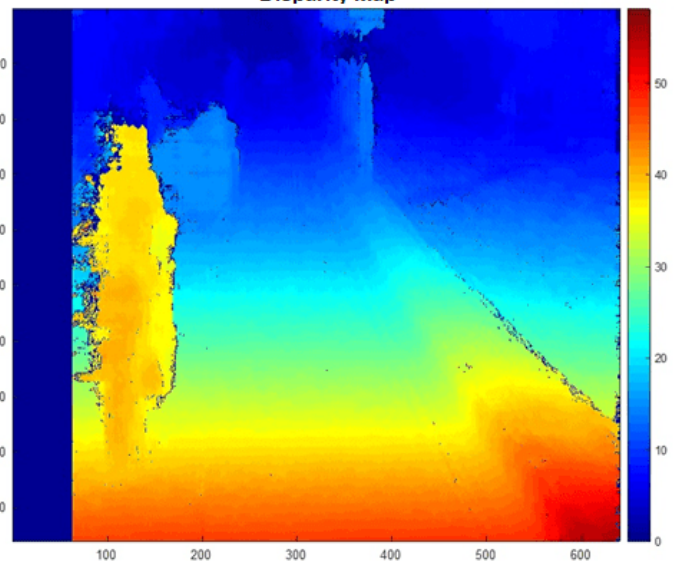
The example model is compatible to generate HDL code. You must have an HDL Coder™ license to generate HDL code. The design was synthesized for the Intel® Arria® 10 GX (115S2F45I1SG) FPGA. The table below shows resource utilization for three disparity level at different image resolutions. Considering one pair of stereo input images as a frame, the algorithm throughput is estimated by finding the number of clock cycles required for processing the current frame before the arrival of next frame. The core algorithm throughput, without overhead of buffering input and output data, is the maximum operating frequency divided by the minimum cycles required between input frames. For example, for 128 disparity levels and 1280-by-720 image resolution, the minimum cycles to process the input frame is 938,857 clock cycles/frame. The maximum operating frequency obtained for algorithm with 128 disparity levels is 61.69 MHz, the core algorithm throughput is computed as 65 frames per second.

```
% =================================================================================
% |Disparity Levels        ||       64       ||       96       ||        128       |
% =================================================================================
% |Input Image Resolution  ||   640 x 480    ||   960 x 540    ||    1280 x 720    |
% |ALM Utilization         ||   45,613 (11%) ||   64,225 (15%) ||    85,194 (20%)  |
% |Total Registers         ||   49,232       ||   64,361       ||    85,564        |
% |Total Block Memory Bits || 3,137,312 (6%) || 4,599,744 (9%) || 11,527,360 (21%) |
% |Total RAM Blocks        ||   264 (10%)    ||   409 (16%)    ||    741 (28%)     |
% |Total DSP Blocks        ||   65 (4%)      ||   97 (6%)      ||    129 (8%)      |
% =================================================================================
```

References

[1] Hirschmuller H., Accurate and Efficient Stereo Processing by Semi-Global Matching and Mutual Information, International Conference on Computer Vision and Pattern Recognition, 2005.

[2] Spangenberg R., Langner T., and Rojas R., Weighted Semi-Global matching and Center-Symmetric Census Transform for Robust Driver Assistance, Computer Analysis of Images and Patterns, 2013.

[3] Gehrig S., Eberli F., and Meyer T., A Real-Time Low-Power Stereo Vision Engine Using Semi-Global Matching, International Conference on Computer Vision System, 2009.

Stereo Image Rectification

This example shows how to implement stereo image rectification for a calibrated stereo camera pair. The example model is FPGA-hardware compatible and provides real-time performance. This example compares its results with the Computer Vision Toolbox™ `rectifyStereoImages` function.

Introduction

A stereo camera is a camera system with two or more lenses with a separate image sensor for each lens. They are used for distance estimation, making 3-D pictures, and stereoviews. Camera lenses distort images, and it is difficult to align two cameras to be perfectly parallel. So, the raw images from a pair of stereo cameras must be rectified. Stereo image rectification projects images onto a common image plane in such a way that the corresponding points in the two stereo images have the same row coordinates. This image projection corrects the images to appear as if the two cameras are parallel.

The algorithm used in this example performs distortion removal and alignment correction in a single system.

Stereo Image Rectification Algorithm

The stereo image rectification algorithm uses a reverse mapping technique to map the pixel locations of the output rectified image to the pixels in the input camera image. The diagram shows the four stages of the algorithm.
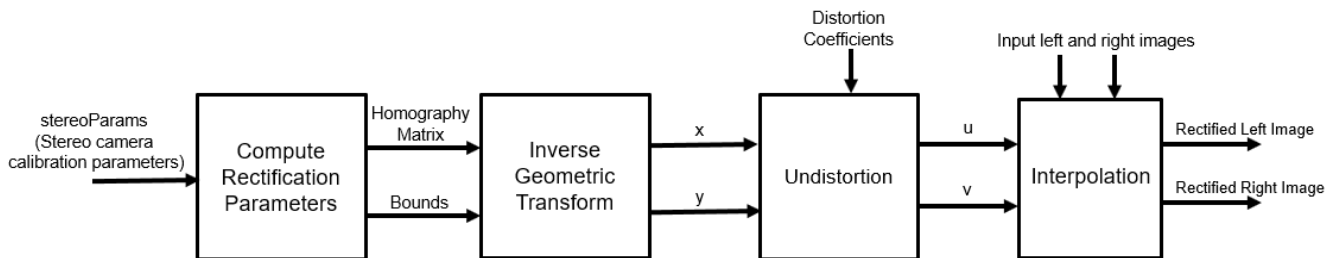
Compute Rectification Parameters: This stage computes rectification parameters from input stereo camera calibration parameters. These calibration parameters include camera intrinsics, rotation matrices, translation matrices, and distortion coefficients (radial and tangential). This stage returns a homography matrix for each camera, and the output bounds. The output bounds are needed to compute the integer pixel coordinates of the output rectified image, and the homography matrices are needed to transform integer pixel coordinates in the output rectified image to corresponding coordinates of the undistorted image.

Inverse Geometric Transform: An inverse geometric transformation translates a point in one image plane onto another image plane. In stereo image rectification, this operation maps integer pixel coordinates in the output rectified image to the corresponding coordinates of the input camera image by using the homography matrix, *H*. If *(p,q)* is an integer pixel coordinate in the rectified output image and *(x,y)* is the corresponding coordinate of the undistorted image, then this equation describes the transformation.

$$[x \ \ y \ \ z]_{1X3} = [p \ \ q \ \ 1]_{1X3} * H^{-1}_{3X3}$$

where H is the homography matrix. To convert from homogeneous to cartesian coordinates, x is set to x/z and y is set to y/z.

Undistortion: Lens distortions are optical aberrations which may deform the images. There are two main types of lens distortions: radial and tangential distortions. Radial distortion occurs when light rays bend more near the edges of a lens than they do at its optical center. Tangential distortion occurs when the lens and the image plane are not parallel. For distortion removal, the algorithm maps the coordinates of the undistorted image to the input camera image by using distortion coefficients.

Let *(u,v)* be the coordinates of the input camera image and *(x,y)* be the undistorted pixel locations. x and y are normalized from pixel coordinates by translating to the optical center and dividing by the focal length in pixels. The following equations describe the undistortion operation.

$$u_{radial} = x(1 + k_1 r^2 + k_2 r^4) \; , \; u_{tangential} = 2p_1 xy + p_2(r^2 + 2x^2)$$

$$v_{radial} = y(1 + k_1 r^2 + k_2 r^4) \; , \; v_{tangential} = 2p_2 xy + p_1(r^2 + 2y^2)$$

where $r^2 = x^2 + y^2$.

k_1, k_2 are *radial distortion coefficients* and p_1, p_2 are *tangential distortion coefficients*.

$$u = u_{radial} + u_{tangential}$$

$$v = v_{radial} + v_{tangential}$$

Inverse geometric transformation and undistortion both contribute to an overall mapping between the coordinates of the output undistorted rectified image *(u,v)* and the coordinates of the input camera image.

Interpolation: Interpolation resamples the image intensity values corresponding to the generated coordinates. The example uses bilinear interpolation.
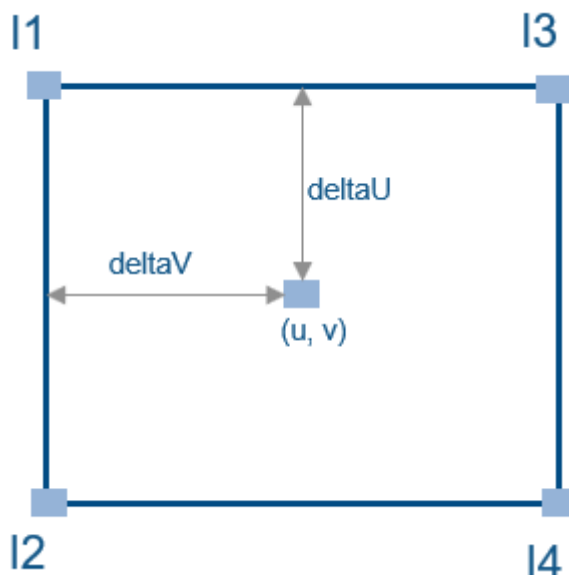
As shown in the diagram, *(u,v)* is the coordinate of the input pixel generated by the undistortion stage. *I1*, *I2*, *I3*, and *I4* are the four neighboring pixels, and *deltaU* and *deltaV* are the displacements of the target pixel from its neighboring pixels. This stage computes the weighted average of the four neighboring pixels by using this equation.

$$rectifiedPixel = I_1(1 - deltaU)(1 - deltaV) + I_2(deltaU)(1 - deltaV) + I_3(1 - deltaU)(deltaV) + I_4(deltaU)(deltaV)$$

HDL Implementation

The figure shows the top-level view of the StereoImageRectificationHDL model. The LeftInputImage and RightInputImage blocks import the stereo left and right images from files. The Frame To Pixels blocks convert these stereo image frames to pixel streams with `pixelcontrol` buses for input to the HDLStereoImageRectification subsystem. This subsystem performs the inverse geometric transform, undistortion, and interpolation to generate the rectified output pixel values. The Pixels To Frame blocks convert the streams of output pixels back to frames. The LeftImageViewer and RightImageViewer subsystems display the input frames and their corresponding rectified outputs.
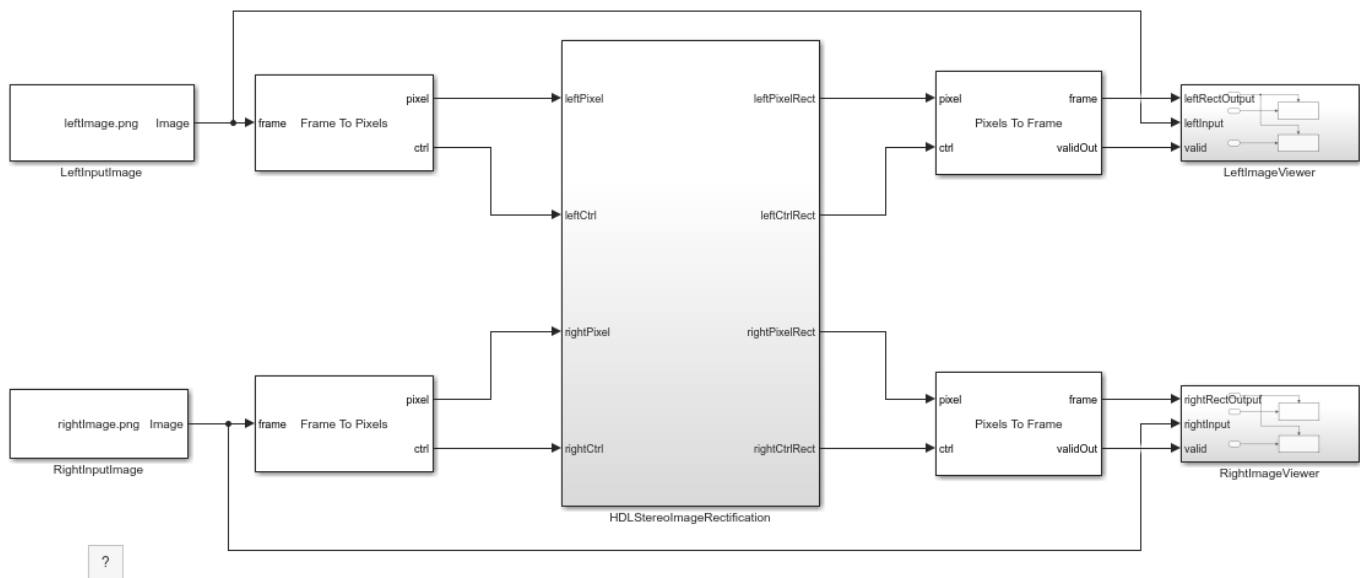
Copyright 2019-2022 The MathWorks, Inc.

The InitFcn of the example model imports the stereo calibration parameters from a data file and computes the rectification parameters by calling `ComputeRectificationParams.m`. Alternatively, you can generate your own set of rectification parameters and provide them as mask parameters of the InverseGeometricTransform and Undistortion subsystems.

The HDLStereoImageRectification subsystem generates a single `pixelcontrol` bus from the two input **ctrl** busses. The RectifiedCoordinateGeneration subsystem generates the row and column pixel coordinates of the output rectified and undistorted image. It uses two HDL counters to generate the row and column coordinates. The InverseGeometricTransform subsystems map these coordinates onto their corresponding row and column coordinates, *(x,y)*, of the distorted image. The Undistortion subsystems map the *(x,y)* coordinates to its corresponding coordinate *(u,v)* of the input camera image, using the distortion coefficients and stereo camera intrinsics.

The Interpolation subsystems store the pixel intensities of the input stereo images in a memory and calculate the addresses of the four neighbors of *(u,v)* required for interpolation. To calculate each

rectified output pixel intensity, the subsystem reads the four neighbor pixel values and finds their weighted sum.
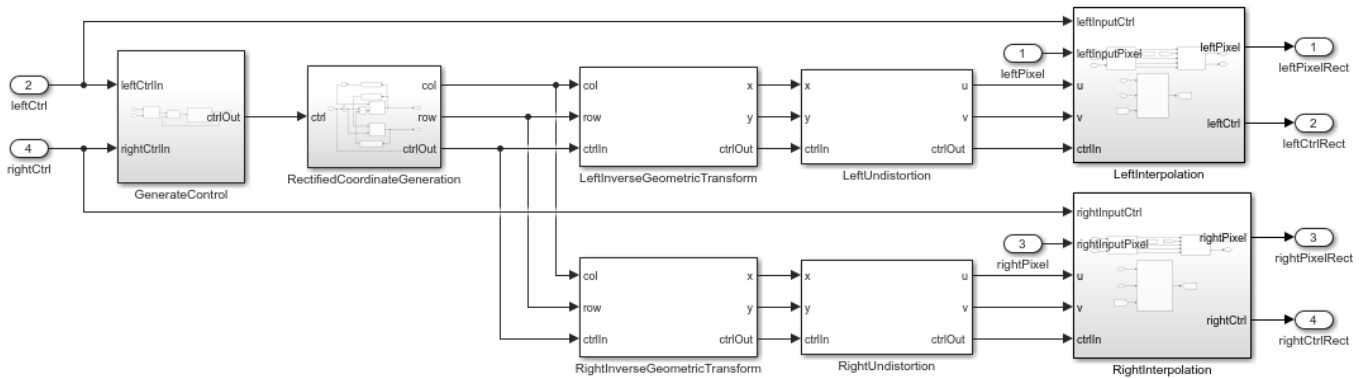
Inverse Geometric Transformation

The HDL implementation of inverse geometric transformation multiplies the coordinates *[row col 1]* with the inverse homography matrix. The inverse homography matrix (3-by-3) is a masked parameter of the InverseGeometricTransformation subsystem. `ComputeRectificationParams.m`, called in the InitFcn of the model, generates the homography matrix. The Transformation subsystem implements the matrix multiplication with Product blocks that multiply by each element of the homography matrix. The HomogeneousToCartesian subsystem converts the generated homogeneous coordinates, *[x y z]* back to the cartesian format, *[x y]* for further processing. The HomogeneousToCartesian subsystem uses a Reciprocal block configured to use the `ShiftAdd` architecture, and the **UsePipelines** parameter is set to `'on'`. To see these parameters, right-click the block and select **HDL Code > HDL Block Properties**. Until this stage, the word length was allowed to grow with each operation. After the HomogeneousToCartesian subsystem, the word length of the coordinates is truncated to a size that still ensures precision and accuracy of the generated coordinates.
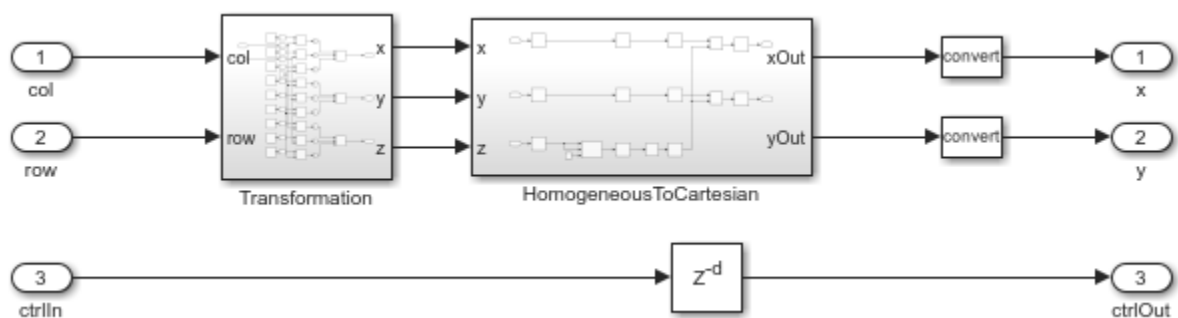
Undistortion

The HDL implementation of Undistortion takes the 3-by-3 camera intrinsic matrix, distortion coefficients *[k1 k2 p1 p2]*, and the reciprocal of *fx* and *fy* as masked parameters. `ComputeRectificationParams.m`, which is called in the InitFcn of the model, generates these parameters. The intrinsic matrix is defined as $\begin{bmatrix} fx & skew & cx \\ 0 & fy & cy \\ 0 & 0 & 1 \end{bmatrix}$.

The Undistortion subsystem implements the equations mentioned in the Stereo Image Rectification Algorithm section by using Sum, Product, and Shift arithmetic blocks. The word length is allowed to grow with each operation, and then the Denormalization subsystem truncates the word length to a size that still ensures the precision and accuracy of the generated coordinates.
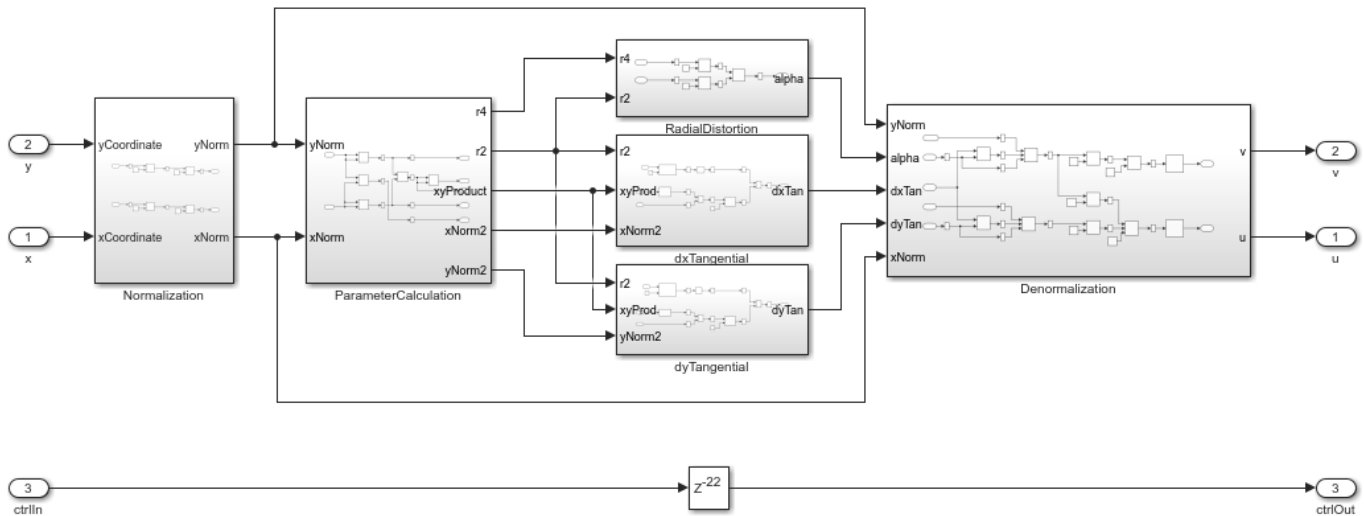
Interpolation

These sections describe the three components inside the Interpolation subsystem.
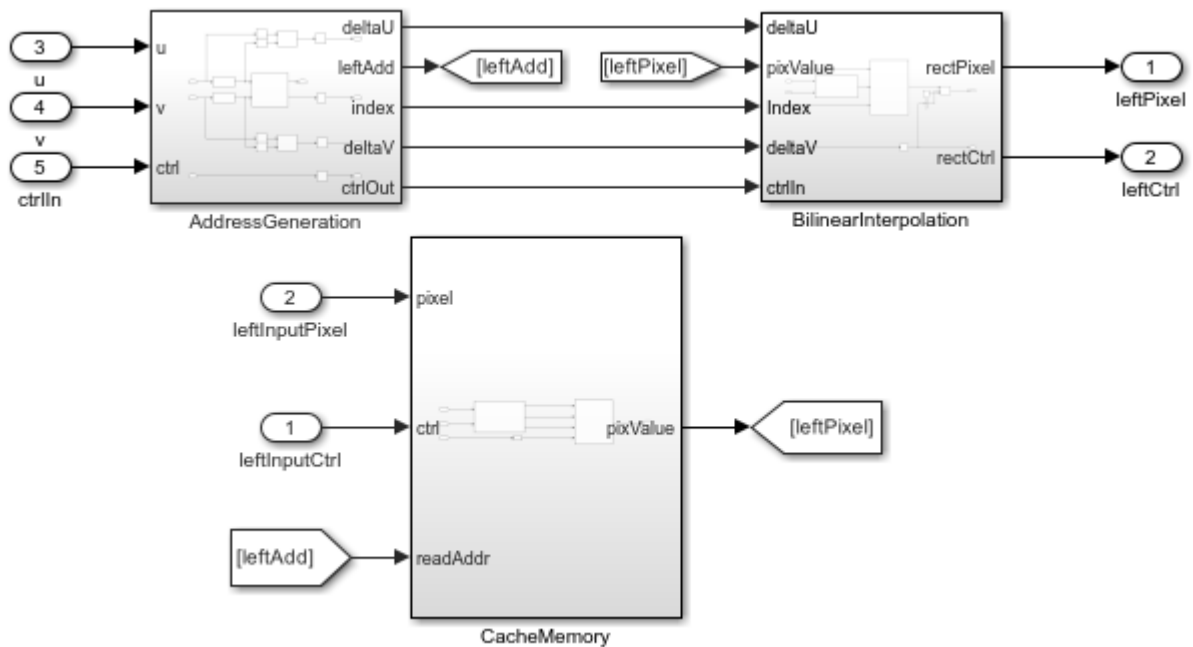
Address Generation

The AddressGeneration subsystem takes the mapped coordinate of the input raw image *(u,v)* as input. It calculates the displacement *deltaU* and *deltaV* of each pixel from its neighboring pixels. It also rounds the coordinates to the nearest integer toward negative infinity.

The AddressCalculation subsystem checks the coordinates against the bounds of the input images. If any coordinate is outside the image dimensions, is capped to the boundary value for further processing. Next, the subsystem calculates the index of the address of each of the four neighborhood pixels in the CacheMemory subsystem. The index represents the column of the cache. The index for each address is determined using the even and odd nature of the incoming column and row coordinates, as determined by the Extract Bits block.

```
% ==========================
% |Row   || Col   || Index ||
% ==========================
% |Odd   || Odd   ||   1   ||
% |Even  || Odd   ||   2   ||
% |Odd   || Even  ||   3   ||
% |Even  || Even  ||   4   ||
% ==========================
```

The address of the neighborhood pixels is generated using this equation:

$$Address = (\frac{Size\,of\,column}{2} * nR) + nC$$

where *nR* is the row coordinate, and *nC* is the column coordinate.

$$nR = \frac{row}{2} - 1 \text{, if row is even }, nR = \frac{row-1}{2} \text{, if row is odd}$$

$$nC = \frac{col}{2} \text{, if col is even }, nC = \frac{col+1}{2} \text{, if col is odd}$$

Once all the addresses and their corresponding indices are generated, they are vectorized using a Vector Concatenate block. The IndexChangeForMemoryAccess MATLAB Function block rearranges the addresses in increasing order of their indices. This operation ensures the correct fetching of the data from the CacheMemory block. The addresses are then given as an input to the CacheMemory block, and the *index*, *deltaU*, and *deltaV* are passed to the BilinearInterpolation subsystem.
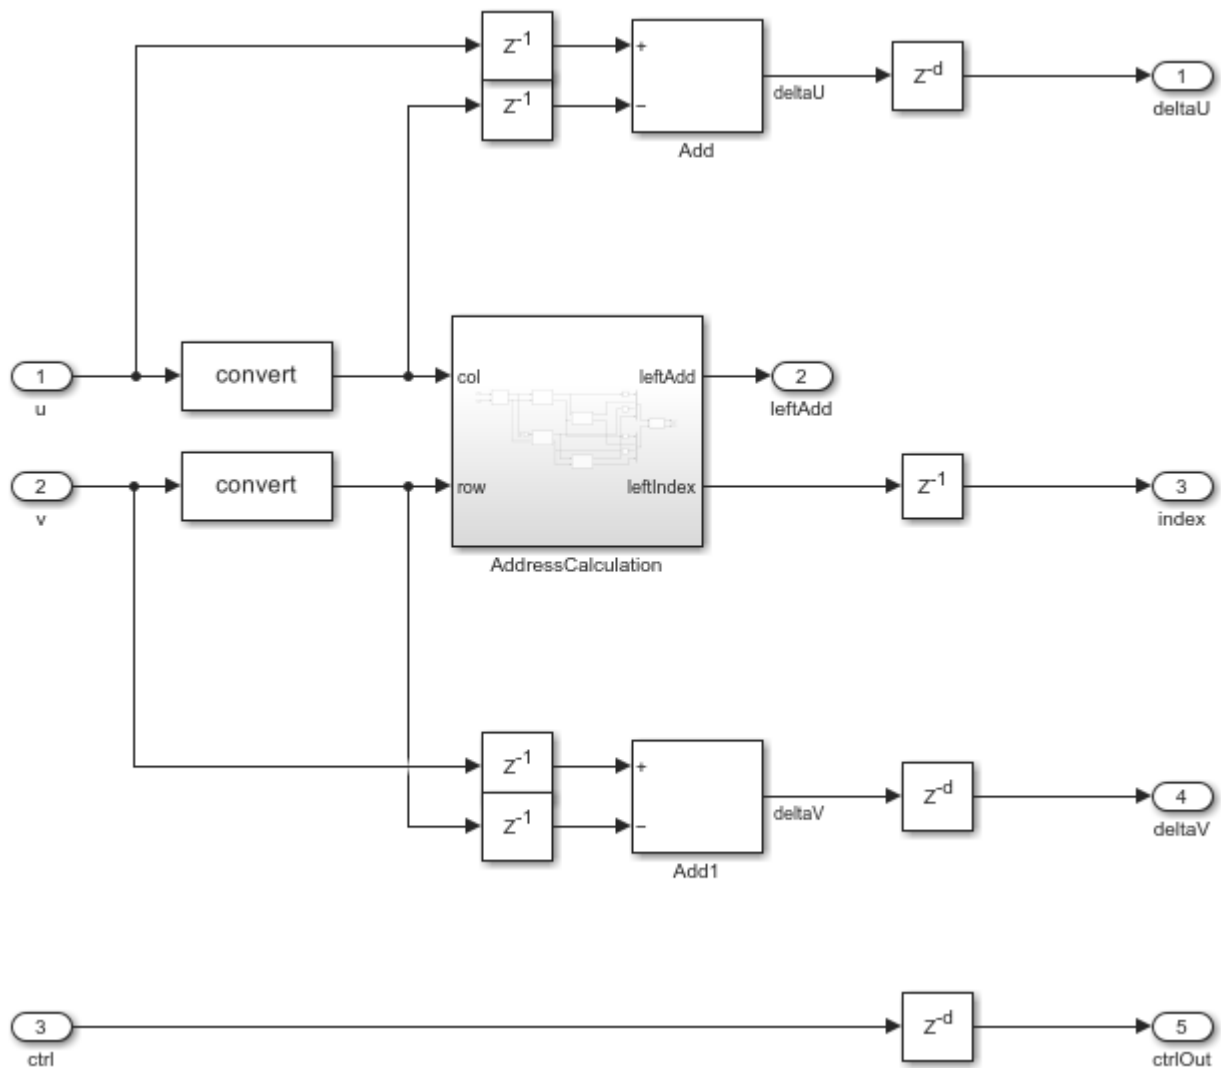
Cache Memory

The CacheMemory subsystem contains a Simple Dual Port RAM block. The input pixels are buffered to form [Line 1 Pixel 1 | Line 2 Pixel 1 | Line 1 Pixel 2 | Line 2 Pixel 2] in the RAM. This configuration enables the algorithm to read all four neighboring pixels in one cycle. The required size of the cache memory is calculated from the *offset* and *displacement* parameters in `ComputeRectificationParams.m` script. The *displacement* is the sum of *maximum deviation* and the *first row map*. The *first row map* is the maximum value of the input image row coordinate that corresponds to the first row of the output rectified image. *Maximum deviation* is the greatest difference between the maximum and minimum row coordinates for each row of the input image row map.

The WriteControl subsystem forms vectors of incoming pixels, and vectors of write enables and write addresses. The AddressGeneration subsystem provides a vector of read addresses. The vector of pixels returned from the RAM are passed to the BilinearInterpolation subsystem.
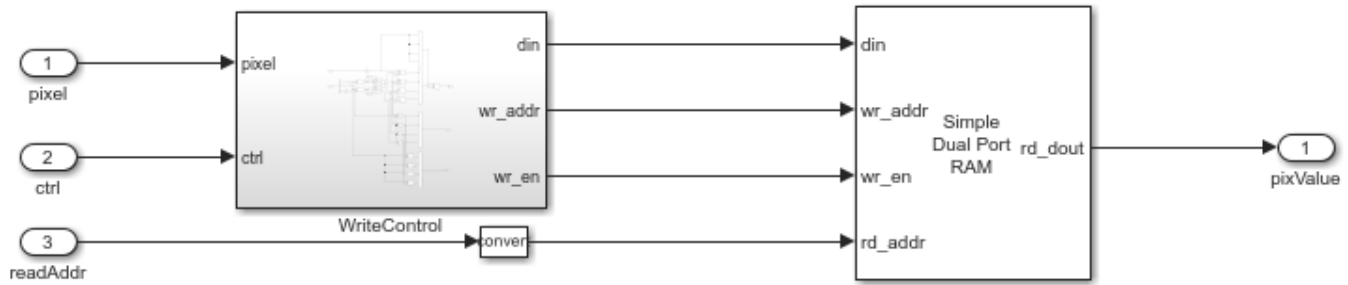
Bilinear Interpolation

The BilinearInterpolation subsystem rearranges the vector of read pixels from the cache to their original indices. Then, the BilinearInterpolationEquation block calculates a weighted sum of the neighborhood pixels by using the bilinear interpolation equation mentioned in the Stereo Image Rectification Algorithm section. The result of the interpolation is the value of the output rectified pixel.
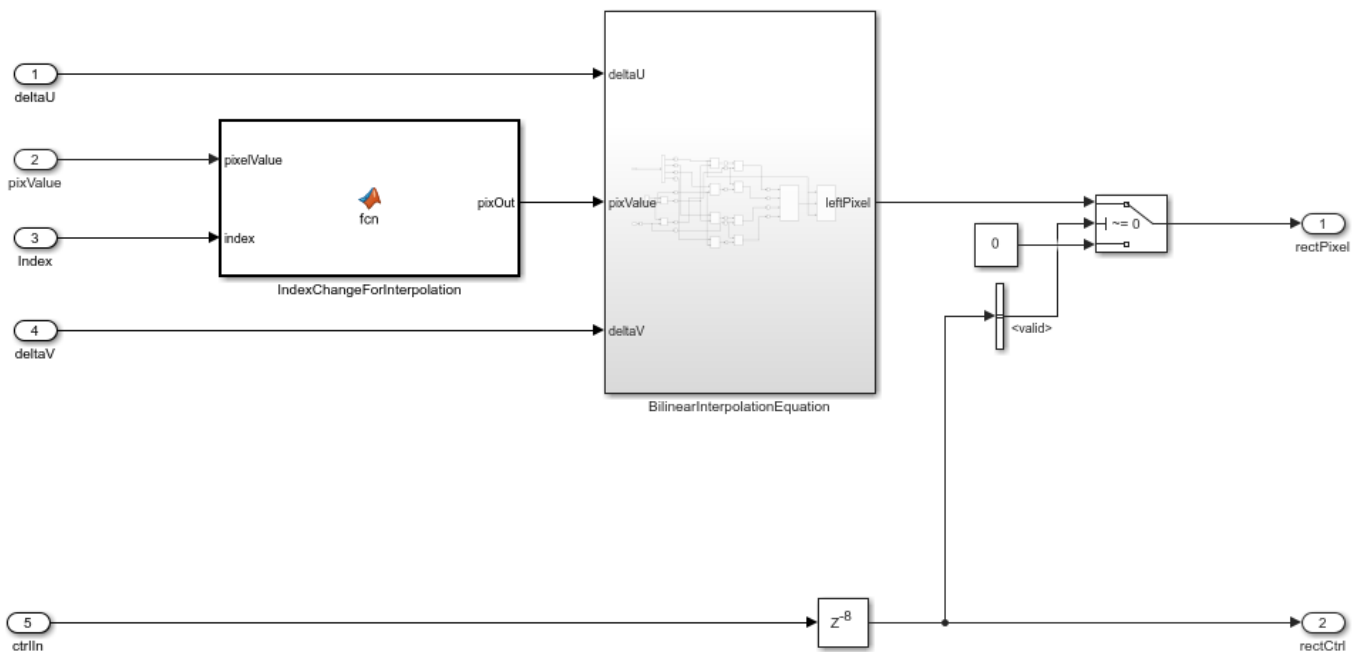
Simulation and Results

This example uses 960-by-1280 stereo images. The input pixels use the `uint8` data type. The example does not provide multipixel support. Due to the large frame sizes used in this example, simulation can take a relatively long time to complete.

The figure shows the left and right input images and the corresponding rectified output images. The results of the StereoImageRectificationHDL model match the output of the `rectifyStereoImages` function in MATLAB with an error of +/-1.
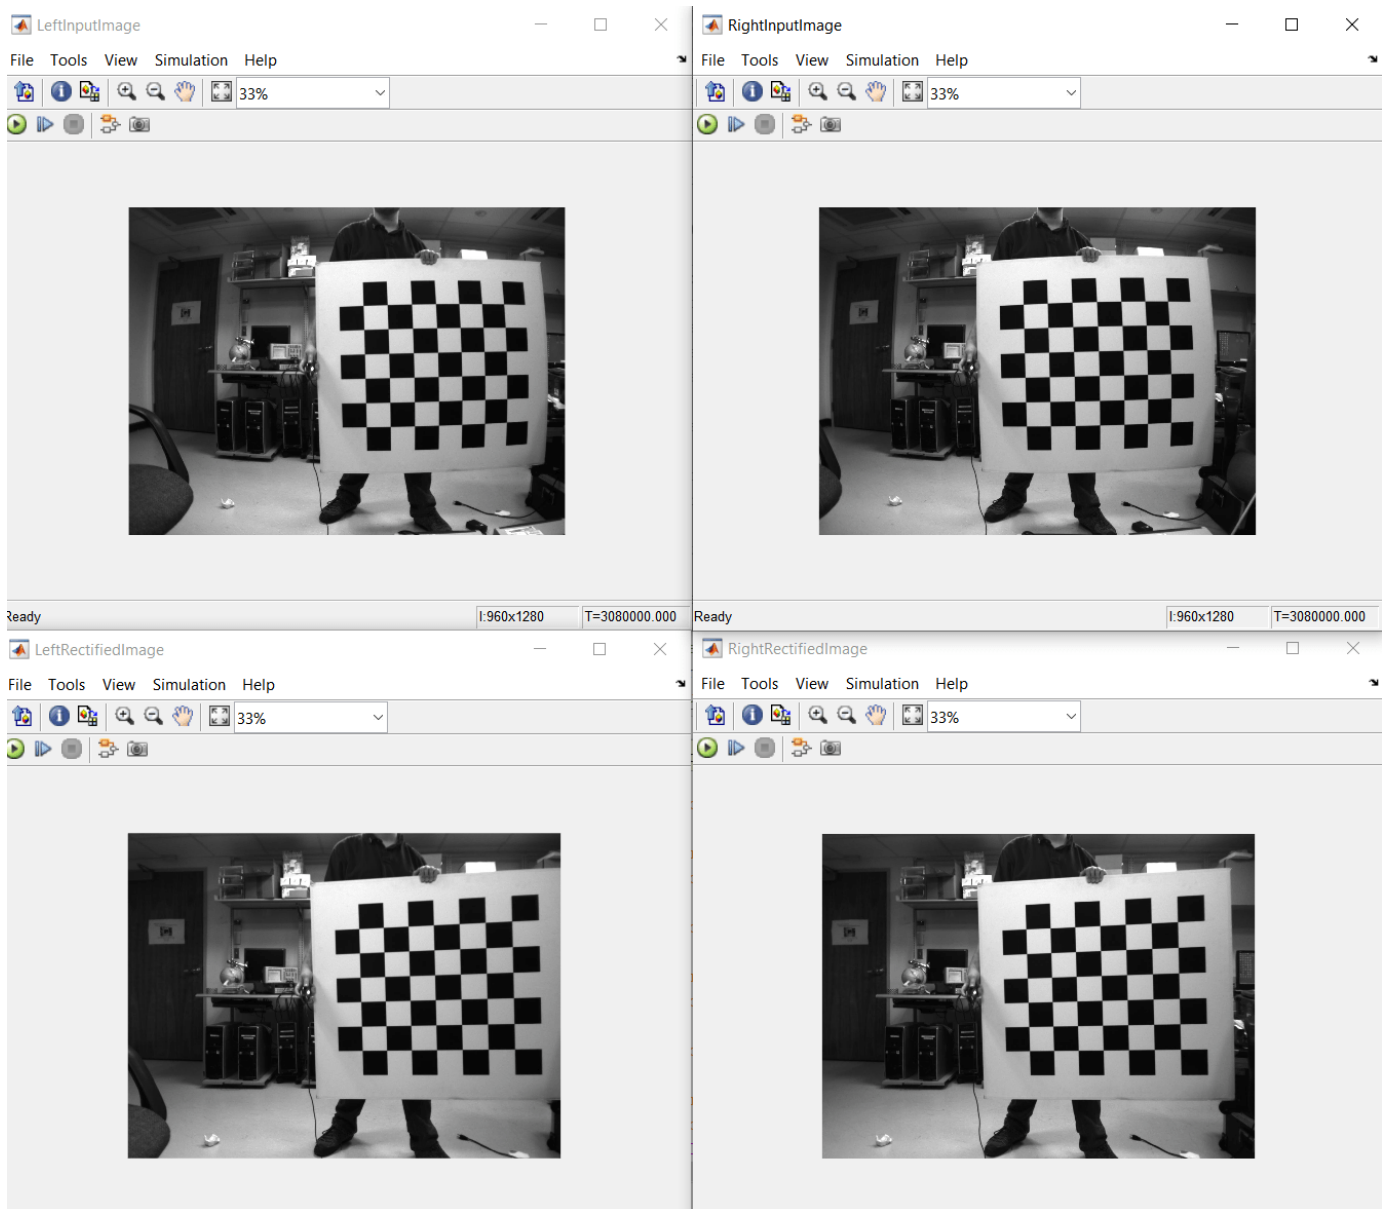
You can generate HDL code for the HDLStereoImageRectification subsystem. You must have an HDL Coder™ license to generate HDL code. This design was synthesized for the Intel® Arria® 10 GX (115S2F45I1SG) FPGA. The HDL design achieves a clock rate of over 150 MHz. The table shows the resource utilization for the subsystem.

```
% ================================================================
% |Model Name              || StereoImageRectificationHDL  ||
% ================================================================
% |Input Image Resolution  ||        960 x 1280            ||
% |ALM Utilization         ||          10675               ||
% |Total Registers         ||          24487               ||
% |Total RAM Blocks         ||          327                ||
% |Total DSP Blocks         ||          218                ||
% ================================================================
```

References

[1] G. Bradski and A. Kaehler, Learning OpenCV : Computer Vision with the OpenCV Library. Sebastopol, CA: O'Reilly, 2008.

See Also

`rectifyStereoImages`

More About

*   "Image Undistortion" on page 2-136
*   "Image Warp" on page 2-147

Image Undistortion

This example shows how to remove lens distortion in images. The algorithm shown here is suitable for FPGAs.

Lens distortions are optical aberrations that deform images. Images typically have two main types of lens distortion: radial or tangential.

Radial distortion occurs when light rays bend more near the edges of a lens than they do at its optical center.
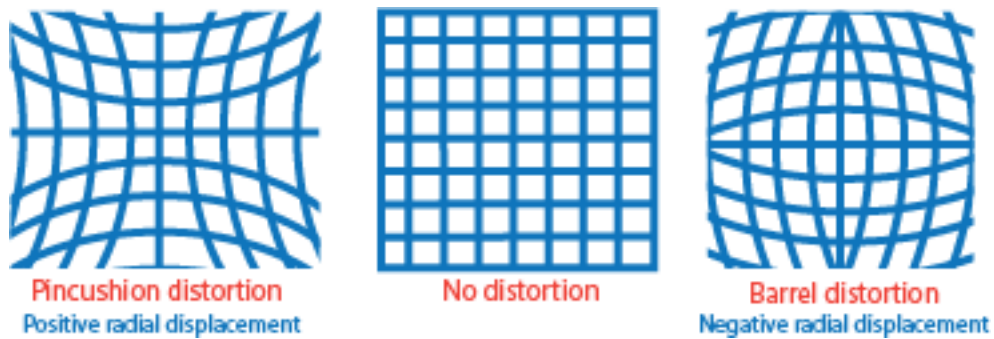
Tangential distortion occurs when the lens and the image plane are not parallel.
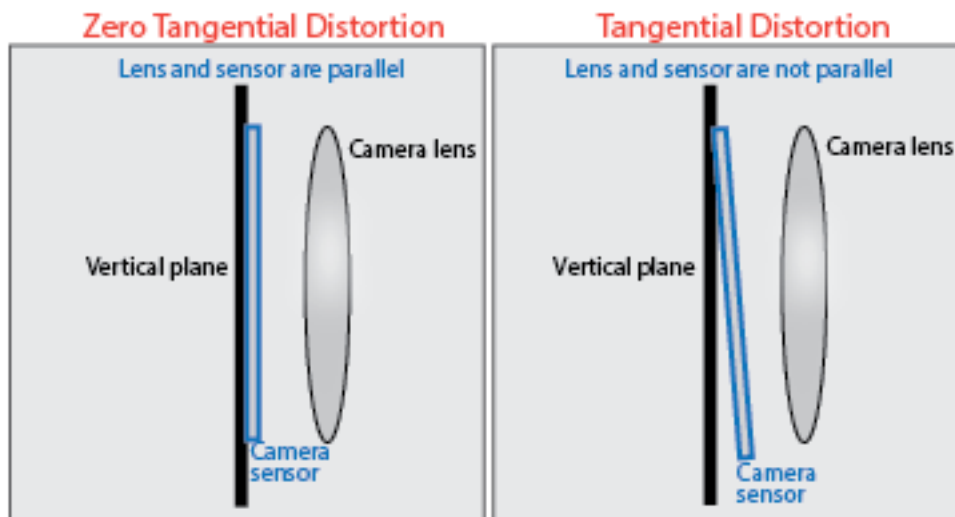
An undistort algorithm maps the coordinates of the output undistorted image to the input camera image by using distortion coefficients. The hardware-friendly undistort implementation in this example performs the same operation as the `imrotate` (Image Processing Toolbox) function.

As inputs to the undistort algorithm, you specify the intrinsic matrix and distortion coefficients that describe the image distortion to be corrected. The intrinsic matrix comprises the focal length, the optical center (also known as the principal point), and the skew coefficient. The distortion coefficients model radial and tangential distortions mathematically.

The `computeCameraParameters` function, included with this example, calculates the input parameters from the specified output image dimensions and a `cameraParameters` (Computer Vision

Toolbox) object that describes the camera intrinsic matrix, distortion coefficients, and camera focal lengths in *x*- and *y*-directions. The `cameraParameters` object is provided in the `cameraParams.mat` file. This function also returns the displacement and offset parameters, which determine how much memory the undistort operation requires.

The example model calls the `computeCameraParameters` function in the PostLoadFcn callback, and then the model removes radial and tangential lens distortions in the input image by using the calculated parameters.

After undistortion, the hardware algorithm calculates the output pixel intensities by using bilinear interpolation. The implementation in this example does not require external DDR memory and instead stores and resamples the output pixel intensities by using on-chip block RAM.

Image Undistortion Algorithm

The image undistortion algorithm maps the pixel locations of the output undistorted image to the pixels in the input distorted image by using a reverse mapping technique. This diagram shows the stages of the algorithm.
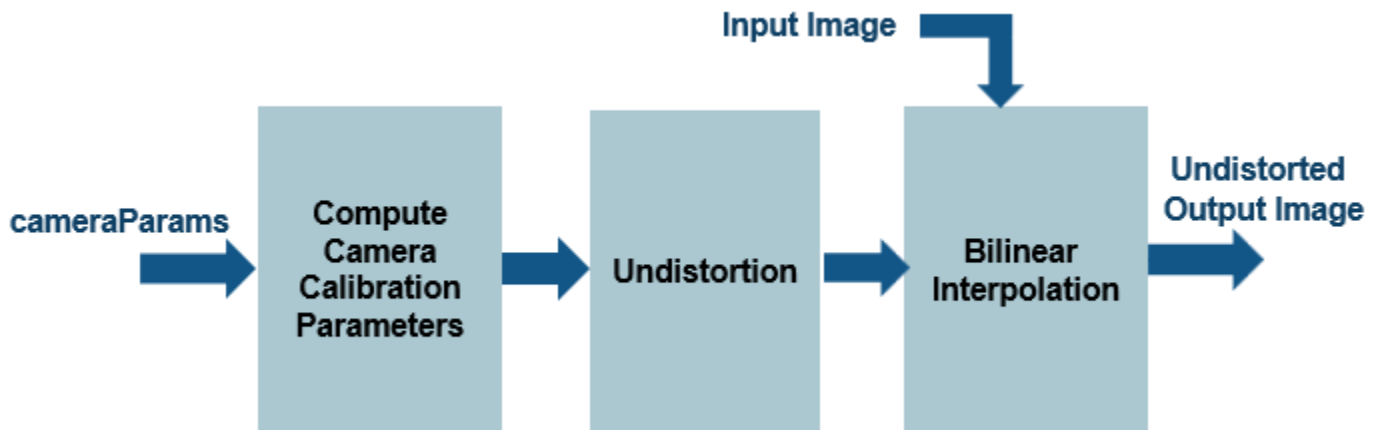
Compute Camera Calibration Parameters

Camera calibration estimates the parameters of a lens and image sensor of an image or video camera. These parameters can be used to correct lens distortion, measure the size of an object in world units, or determine the location of the camera in the scene. Applications such as machine vision, robotic navigation systems, and 3-D scene reconstruction use these operations to detect and measure objects. Camera parameters include intrinsics, extrinsics, and distortion coefficients. This stage computes these parameters from a given input `cameraParameters` object and feeds them to the undistortion stage.

Undistortion

This stage removes distortion by using the distortion coefficients and the intrinsic matrix. These equations model distortion removal.

Let *(u,v)* be the coordinates of the input camera image and *(x,y)* be the undistorted pixel locations. Normalize *x* and *y* from pixel coordinates by translating to the optical center and dividing by the focal length in pixels.

$$u_{radial} = x(1 + k_1 r^2 + k_2 r^4),$$

$$u_{tangential} = 2p_1xy + p_2(r^2 + 2x^2),$$

$$v_{radial} = y(1 + k_1r^2 + k_2r^4),$$

$$v_{tangential} = 2p_2xy + p_1(r^2 + 2y^2),$$

where $r^2 = x^2 + y^2$.

k_1, k_2 are *radial distortion coefficients*, and p_1, p_2 are *tangential distortion coefficients*. Then, calculate the final coeffient values by combining the radial and tangential components.

$$u = u_{radial} + u_{tangential}$$

$$v = v_{radial} + v_{tangential}$$

Undistortion defines a mapping between the coordinates of the output undistorted image *(u,v)* and the coordinates of the distorted input camera image, *(x,y)*.

Bilinear Interpolation

The image undistortion algorithm can produce noninteger values of *(u,v)*. Generating the pixel intensity at each integer position requires a resampling technique such as interpolation. This example resamples the image intensity values corresponding to the generated coordinates by using bilinear interpolation.

In the equation and the diagram, *(u,v)* is the coordinate of the input pixel generated by the undistortion stage. I_1, I_2, I_3, and I_4 are the four neighboring pixels, and ΔU and ΔV are the displacements of the target pixel from its neighboring pixels. This stage of the algorithm computes the weighted average of the four neighboring pixels by using this equation.

$$outputPixel = I_1(1 - \Delta U)(1 - \Delta V) + I_2(\Delta U)(1 - \Delta V) + I_3(1 - \Delta U)(\Delta V) + I_4(\Delta U)(\Delta V)$$
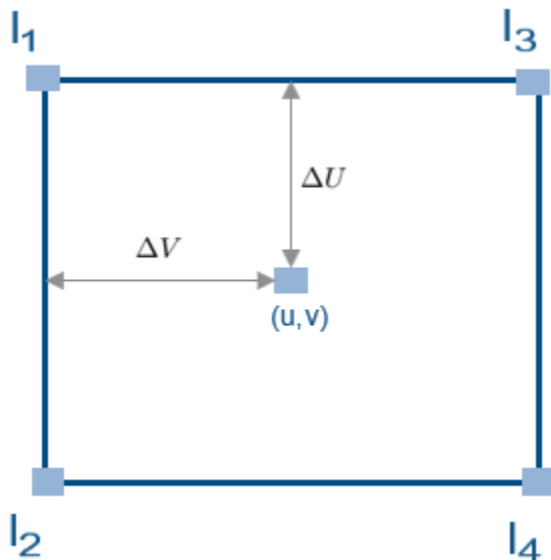
HDL Implementation

This figure shows the top-level view of the `ImageUndistortHDL` model. The InputImage block imports the image from a file. The Frame To Pixels block converts the input image frames to a pixel stream and a `pixelcontrol` bus for input to the `ImageUndistortHDLAlgorithm` subsystem. This subsystem removes distortions from the input image by using the distortion coefficients that you specify in the mask parameters. The Pixels To Frame block converts the stream of output pixels to a frame. The `ImageViewer` subsystem displays the input frame and the corresponding undistorted output.

```
open_system('ImageUndistortHDL');
set(allchild(0),'Visible','off');
```
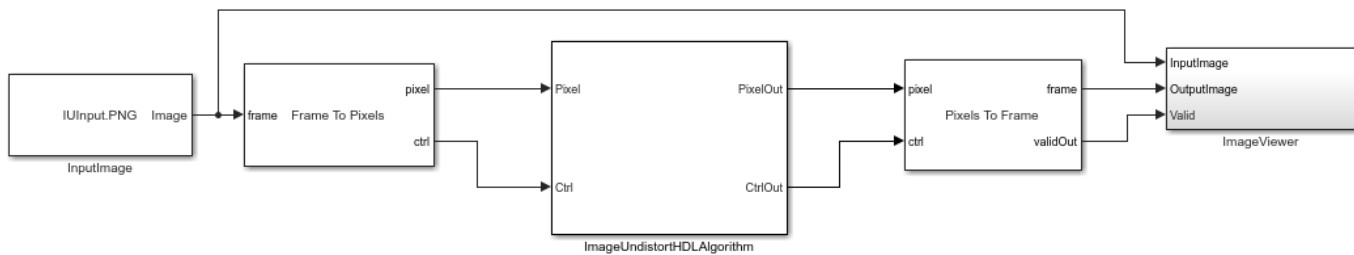

Copyright 2021 The MathWorks, Inc.

The PostLoadFcn callback of the example model imports the camera calibration parameters from the `cameraParams.mat` data file and computes the calibration parameters by calling the `ComputeCameraParameters` function provided with this example. Alternatively, you can generate your own camera calibration parameters and provide them as mask parameters of the `ImageUndistortHDLAlgorithm` subsystem.

In the `ImageUndistortHDLAlgorithm` subsystem, the `GenerateControl` subsystem uses the displacement parameter to modify the `pixelcontrol` bus from the input `ctrl` port. The `CoordinateGeneration` subsystem generates the row and column pixel coordinates *(x,y)* of the output undistorted image by using two HDL counters. The `Undistortion` subsystem maps the *(x,y)* position to its corresponding *(u,v)* position of the input camera image by using the distortion coefficients and camera intrinsics.

The `AddressGeneration` subsystem calculates the addresses of the four neighbors of *(u,v)* required for interpolation. This subsystem also computes the parameters ΔU, ΔV, $Bound$, and $IndexVector$, required for bilinear interpolation.

The `Interpolation` subsystem stores the pixel intensities of the input image in a memory modeled with a Simple Dual Port RAM block. To calculate each output pixel intensity, the subsystem reads the four neighbor pixel values and computes their weighted sum.

```
open_system('ImageUndistortHDL/ImageUndistortHDLAlgorithm','force');
```
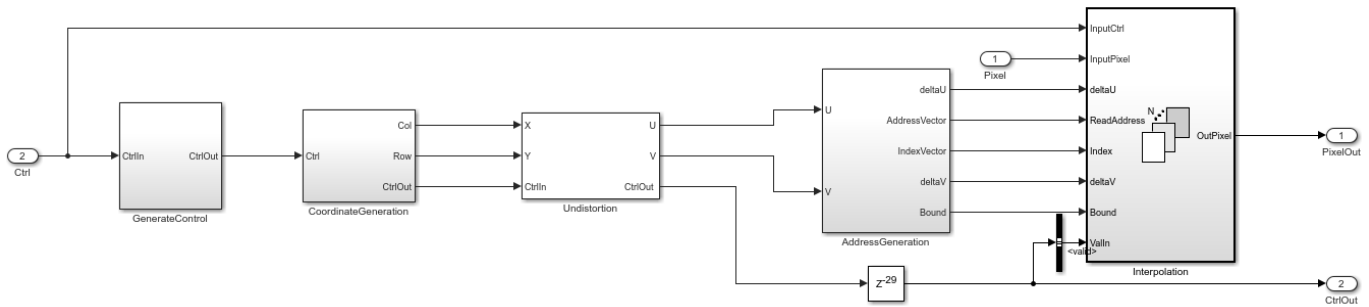
Undistortion

The HDL implementation of undistortion takes the 3-by-3 camera intrinsic matrix, distortion coefficients $[k_1 \, k_2 \, p_1 \, p_2]$, and the reciprocal of f_x and f_y as masked parameters. The `ComputeCameraParameters` function, which is called in the PostLoadFcn callback of the model, generates these parameters. The intrinsic matrix is:

$$\begin{bmatrix} f_x & skew & c_x \\ 0 & f_y & c_y \\ 0 & 0 & 1 \end{bmatrix}$$

The `Undistortion` subsystem implements the equations mentioned in the Image Undistortion Algorithm section by using Sum, Product, and Shift arithmetic blocks. The word length grows with each operation, and then the `Denormalization` subsystem truncates the word length to a size that ensures the precision and accuracy of the generated coordinates.

```
open_system('ImageUndistortHDL/ImageUndistortHDLAlgorithm/Undistortion','force');
```
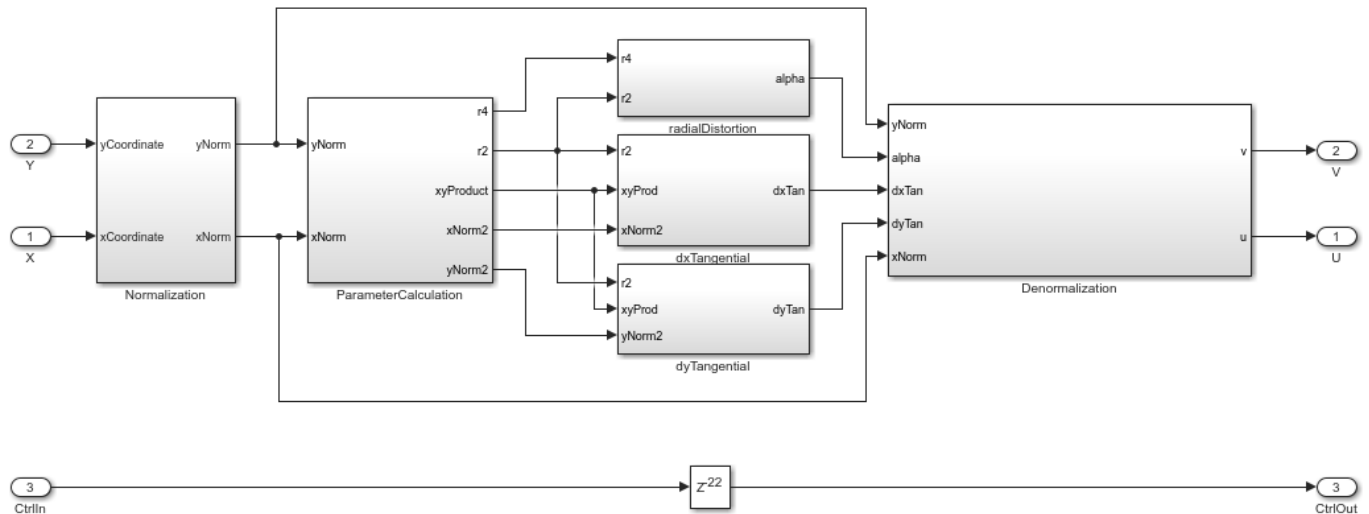

Address Generation

The `AddressGeneration` subsystem calculates the displacement ΔU and ΔV of each pixel from its neighboring pixels by using the mapped coordinate *(u,v)* of the input raw image. The subsystem also rounds the coordinates to the nearest integer toward negative infinity.

```
open_system('ImageUndistortHDL/ImageUndistortHDLAlgorithm/AddressGeneration','force');
```
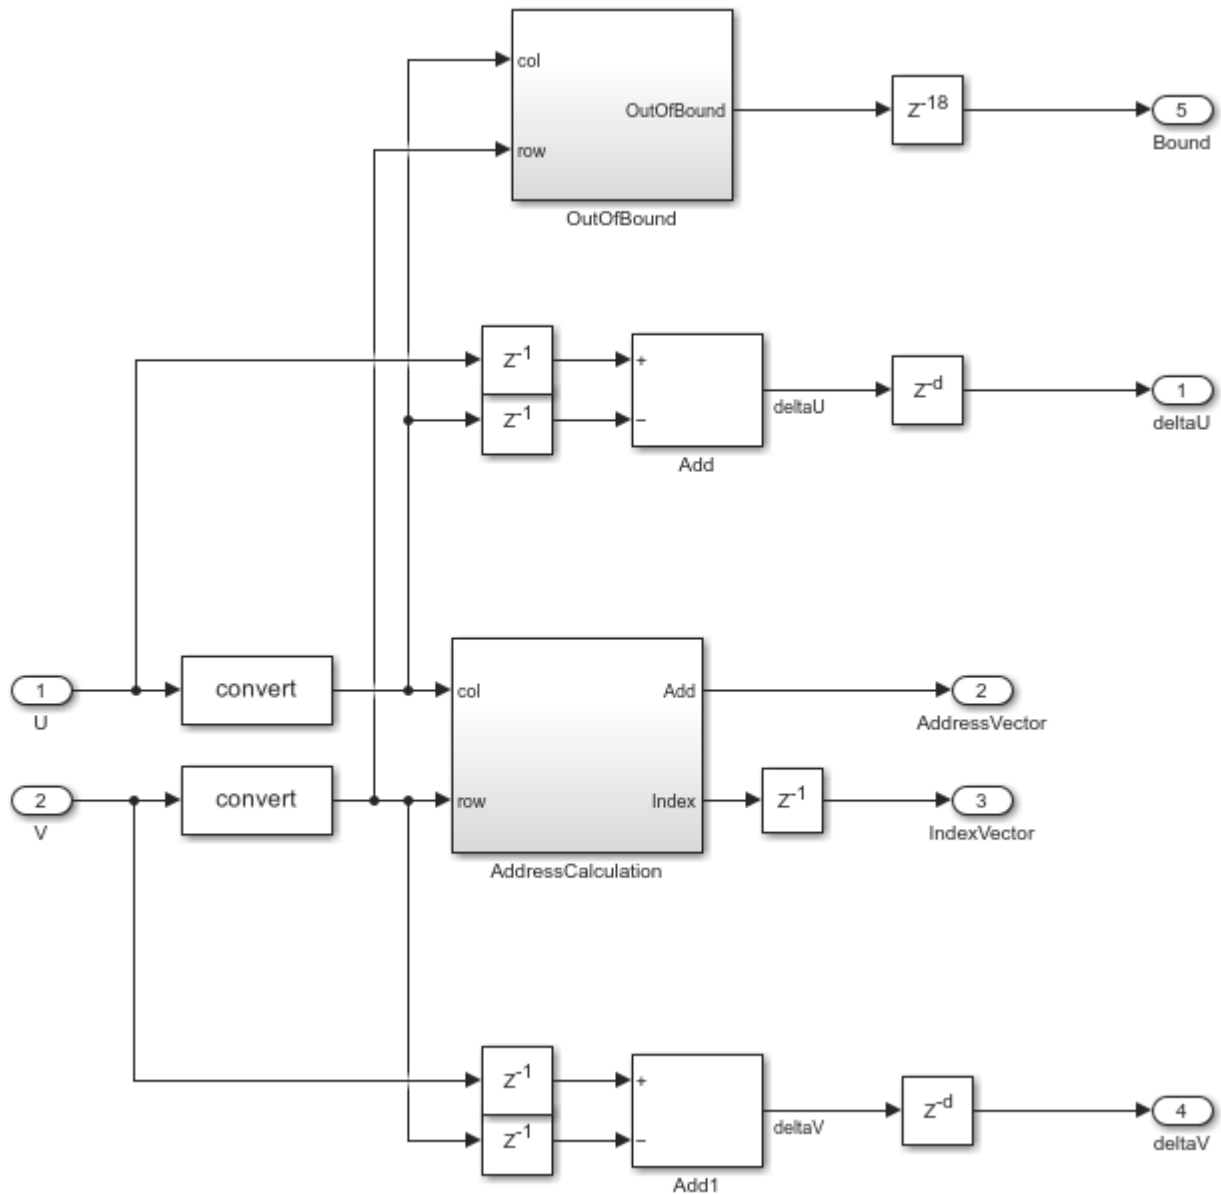

The `AddressCalculation` subsystem checks the coordinates against the bounds of the input images. If any coordinate is outside the image dimensions, the subsystem sets the coordinate to the boundary value. Next, the subsystem calculates the index of the address for each of the four neighborhood pixels in the `CacheMemory`. The index represents the column of the cache. The subsystem finds the index for each address by using the even and odd nature of the incoming column and row coordinates, as determined by the Extract Bits block.

```
% ==========================
% |Row   || Col   || Index ||
% ==========================
% |Odd   || Odd   ||   1   ||
```

```
% |Even || Odd  ||    2    ||
% |Odd  || Even ||    3    ||
% |Even || Even ||    4    ||
% ==========================
```

This equation specifies the address of the neighborhood pixels.

$$addr = (\frac{colSize}{2} * R) + C$$

R is the row coordinate and C is the column coordinate. When `row` is even, then $nR = \frac{row}{2} - 1$. When `row` is odd, then $R = \frac{row-1}{2}$. When `col` is even, then $C = \frac{col}{2}$. When `col` is odd, then $C = \frac{col+1}{2}$.

The IndexChangeForMemoryAccess MATLAB Function block in the `AddressCalculation` subsystem rearranges the addresses in increasing order of their indices. This operation ensures the correct fetching of data from the CacheMemory block. This subsystem passes the addresses to the `CacheMemory` subsystem, and passes `Index`, ΔX, and ΔY to the `Interpolation` subsystem.

The `OutOfBound` subsystem checks whether the *(u,v)* coordinates are out of bounds (that is, if any coordinate is outside the image dimensions). If the coordinate is out of bounds, the subsystem sets the corresponding output pixel to an intensity value of `0`.

Finally, a Vector Concatenate block creates vectors of the addresses and indices.

Interpolation

The `Interpolation` subsystem is a For Each block, which replicates its operation depending on the dimensions of the input pixel. For example, if the input is an RGB image, then the input pixel dimensions are 1-by-3, and the model includes three instances of this operation. Because the model uses a For Each block, it supports RGB or grayscale input. The operation inside the `Interpolation` subsystem comprises two subsystems: `BilinearInterpolation` and `CacheMemory`.

```
open_system('ImageUndistortHDL/ImageUndistortHDLAlgorithm/Interpolation','force');
```

Cache Memory

The `CacheMemory` subsystem contains a Simple Dual Port RAM block. The subsystem buffers the input pixels to form [Line 1 Pixel 1 | Line 2 Pixel 1 | Line 1 Pixel 2 | Line 2 Pixel 2] in the RAM. By using this configuration, the algorithm can read all four neighboring pixels in one cycle. The example calculates the required size of the cache memory from the *offset* output of the `ComputeCameraParams` function. The offset is the sum of the *maximum deviation* and the *first row map*. The *first row map* is the maximum value of the input image row coordinate that corresponds to the first row of the output undistorted image. The *maximum deviation* is the greatest difference between the maximum and minimum row coordinates for each row of the input image row map.

The `WriteControl` subsystem forms vectors of incoming pixels, write enables, and write addresses. The `AddressGeneration` subsystem provides a vector of read addresses. The vector of pixels from the RAM is the input to the `BilinearInterpolation` subsystem.

```
open_system('ImageUndistortHDL/ImageUndistortHDLAlgorithm/Interpolation/CacheMemory','force');
```

Bilinear Interpolation

The `BilinearInterpolation` subsystem rearranges the vector of read pixels from the cache to their original indices. Then, the `BilinearInterpolationEquation` subsystem calculates a weighted sum of the neighborhood pixels by using the bilinear interpolation equation in the Image Undistortion Algorithm section. The result of the interpolation is the value of the output undistorted pixel.

```
open_system('ImageUndistortHDL/ImageUndistortHDLAlgorithm/Interpolation/BilinearInterpolation','
```
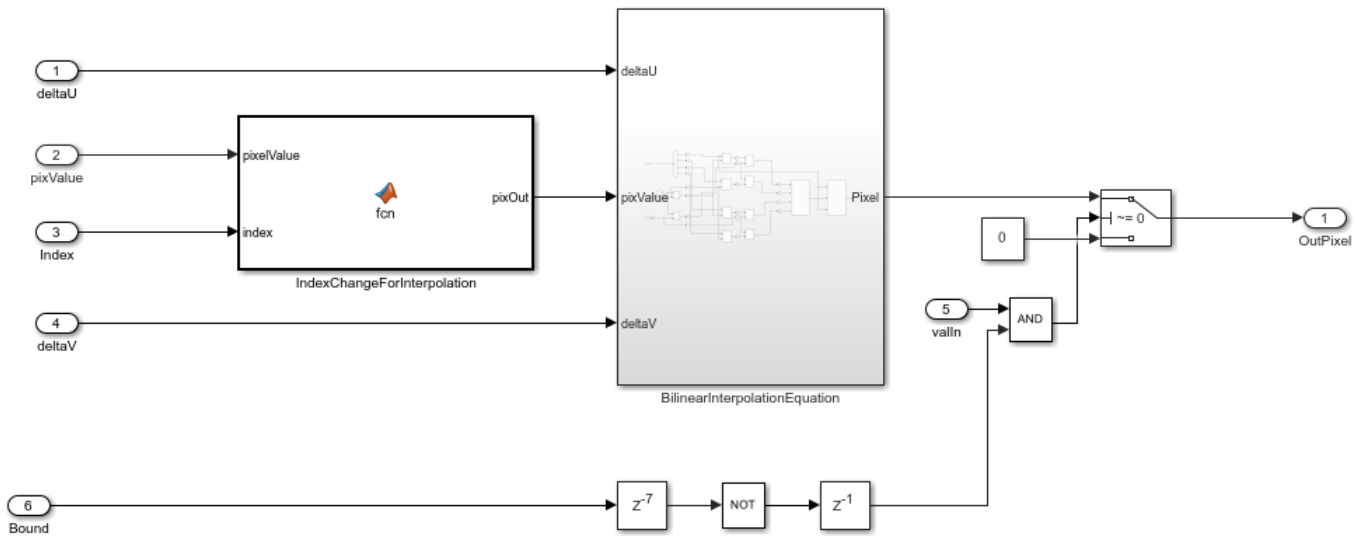

Simulation and Results

This example uses a 510-by-510 grayscale input image. The input pixels use the `uint8` data type for either grayscale or RGB input images.

This figure shows the input distorted image and the corresponding output undistorted image for the camera parameters provided in `cameraParams.mat`. The results of the `ImageUndistortHDL` model for this input matches the output of the `undistortImage` function.
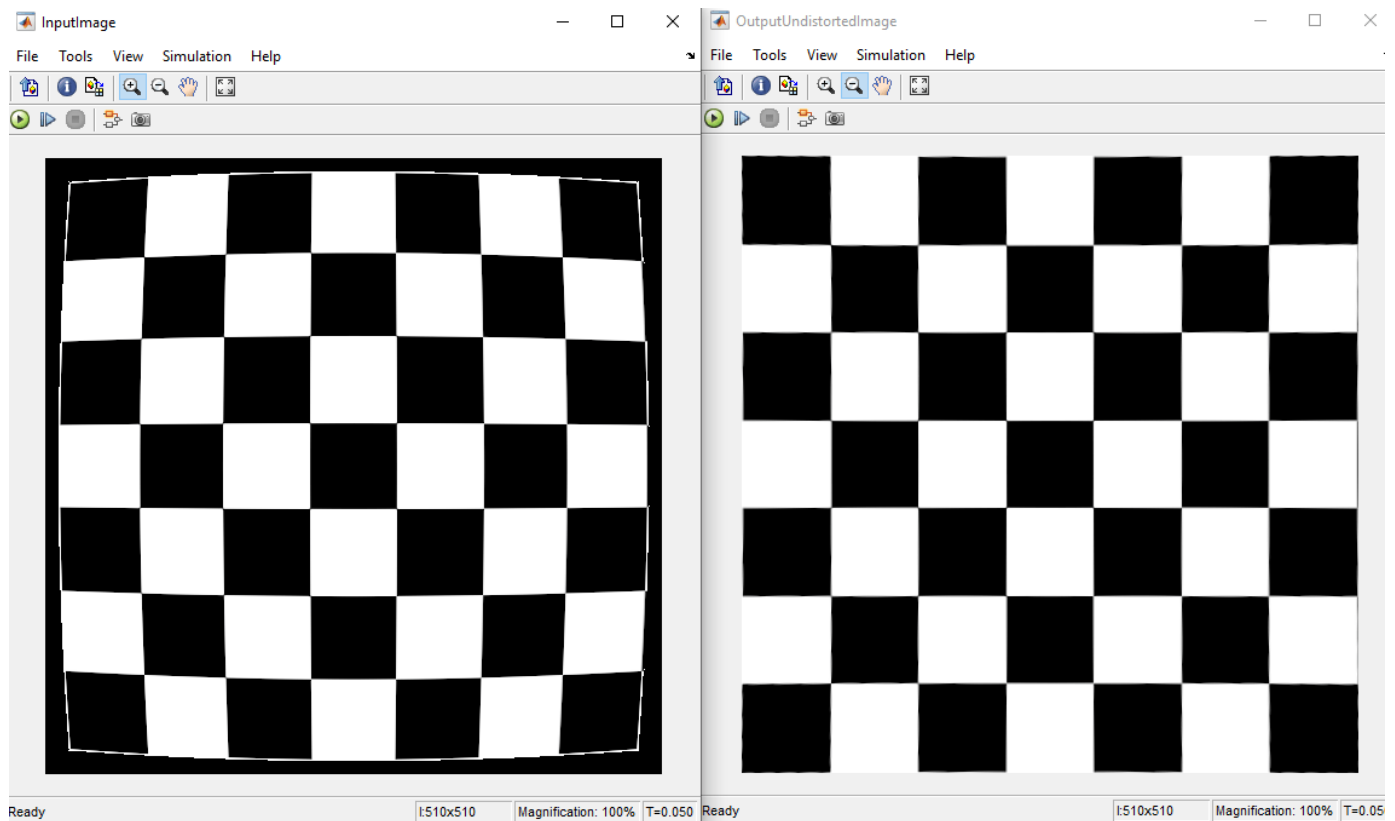
To check and generate the HDL code referenced in this example, you must have the HDL Coder™ product.

To generate the HDL code, use this command.

```
makehdl('ImageUndistortHDL/ImageUndistortHDLAlgorithm')
```

To generate the test bench, use this command.

```
makehdltb('ImageUndistortHDL/ImageUndistortHDLAlgorithm')
```

This design was synthesized using Xilinx® Vivado® for the Xilinx (R) Zynq® ZC706 device and met a timing requirement of over 200 MHz. This table shows the resource utilization for the HDL subsystem.

```
% ============================================================
% |Model Name            ||        ImageUndistortHDL      ||
% ============================================================
% |Input Image Resolution||         510 x 510             ||
% |LUT                   ||           2806                ||
% |FF                    ||           2832                ||
% |BRAM                  ||            17                 ||
```

```
% |Total DSP Blocks         ||              53                ||
% ================================================================
```

See Also

More About

- "Stereo Image Rectification" on page 2-126
- "Image Warp" on page 2-147

Image Warp

This example shows how to implement affine and projective transforms for FPGAs.

Image warping is a common technique in image processing and computer graphics. This technique generates an image to specified requirements by geometrically distorting an input image, an approach that is closely related to the morphing technique. Image warping has diverse applications, such as registration in remote sensing and creating visual special effects in the entertainment industry.

The warp algorithm maps locations in the output image to corresponding locations in the input image, a process known as inverse mapping. The hardware-friendly warp implementation in this example performs the same operation as `imwarp` (Image Processing Toolbox) function.

The algorithm in this example performs an inverse geometric transform and calculates the output pixel intensities by using bilinear interpolation. This implementation does not require external DDR memory and instead resamples the output pixel intensities by using the on-chip BRAM memory.

Image Warp Algorithm

The image warping algorithm maps the pixel locations of the output warped image to the pixel locations in the input image by using a reverse mapping technique. This diagram shows the stages of the algorithm.
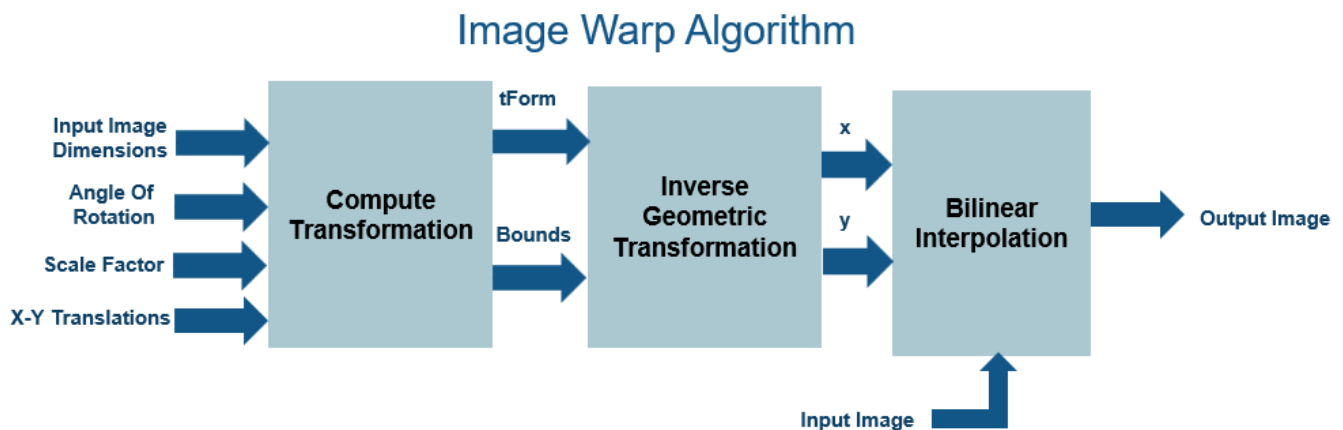
Compute Transformation: This stage computes the inverse transformation matrix. The calculated transformation parameters include the output bounds and the transformation matrix *tForm*. The algorithm requires these bounds to compute the integer pixel coordinates of the output image. The algorithm requires *tForm* to transform the integer pixel coordinates in the output image to the corresponding coordinates of the input image.

Inverse Geometric Transform: An inverse geometric transformation translates a point in one image plane onto another image plane. In image warping, this operation maps the integer pixel coordinates in the output image to the corresponding coordinates of the input image by using the transformation matrix. If *(u,v)* is an integer pixel coordinate in the warped output image and *(x,y)* is the corresponding coordinate of the input image, then this equation describes the transformation.

$$[x \quad y \quad z]_{1-by-3} = [u \quad v \quad 1]_{1-by-3} \cdot tForm_{3-by-3}$$

2-147

tForm is the inverse transformation matrix. To convert from homogeneous to cartesian coordinates, $x = x/z$ and $y = y/z$.

Bilinear Interpolation: The warping algorithm can produce coordinates *(x,y)* with noninteger values. To generate the pixel intensity values at each integer position, a warp algorithm can use various resampling techniques. The example uses bilinear interpolation. Interpolation resamples the image intensity values corresponding to the generated coordinates.

HDL Implementation

The figure shows the top-level view of the ImageWarpHDL model. The Input Image block imports the images from files. The Frame To Pixels block converts the input image frames to a pixel stream and a `pixelcontrol` bus as inputs to the `ImageWarpHDLAlgorithm` subsystem. This subsystem takes these mask parameters.

- **Number of input lines to buffer** — The provided `ComputeImageWarpCacheOffset` function calculates this parameter from the transformation matrix.
- **Input active pixels** — Horizontal size of the input image.
- **Input active lines** — Vertical size of the input image.

The `ImageWarpHDLAlgorithm` subsystem warps the input image as specified by the value of the **tForm** input port. The Pixels To Frame block converts the streams of output pixels back to frames. The `ImageViewer` subsystem displays the input frame and the corresponding warped output.

```
open_system('ImageWarpHDL');
set(allchild(0),'Visible','off');
```
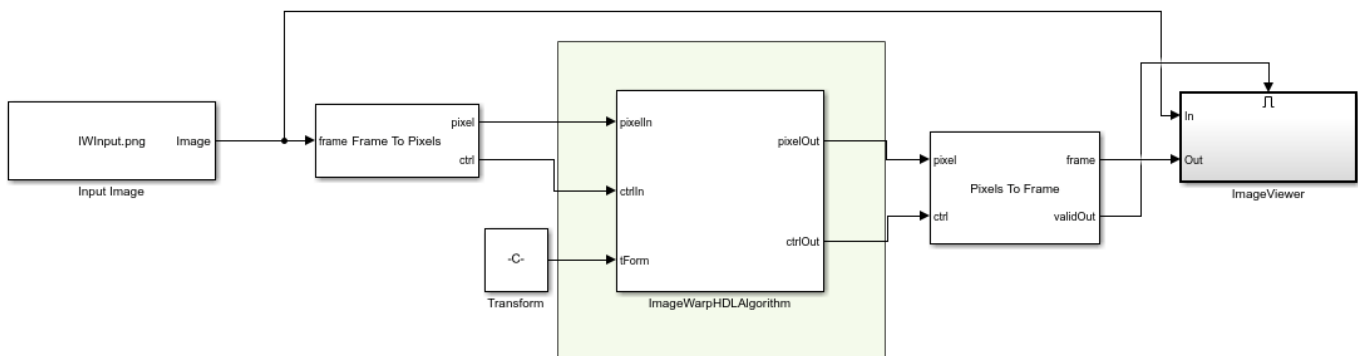

Copyright 2021 The MathWorks, Inc.

The InitFcn callback function loads the transformation matrix from `tForm.mat`. Alternatively, you can generate your own transformation matrix (in the form of a nine-element column vector) and use this vector as the input to the `ImageWarpHDLAlgorithm` subsystem. The InitFcn callback function of the example model also computes the cache offset by calling the `ComputeImageWarpCacheOffset` function. This function calculates the offset and displacement parameters of the output image from the transformation matrix and output image dimensions.

In the `ImageWarpHDLAlgorithm` subsystem, the `GenerateControl` subsystem uses the displacement parameter to modify the `pixelcontrol` bus from the input `ctrl` port. The `CoordinateGeneration` subsystem generates the row and column pixel coordinates *(u,v)* of the

output image by using two HDL counters. The `InverseTransform` subsystem maps these coordinates onto their corresponding coordinates *(x,y)* of the input image.

The `AddressGeneration` subsystem calculates the addresses of the four neighbors of *(x,y)* required for interpolation. This subsystem also computes the parameters $DeltaX$, ΔY, $Bound$, and $indexVector$, which the model uses for bilinear interpolation.

The `Interpolation` subsystem stores the pixel intensities of the input image in a memory. To calculate each output pixel intensity, the subsystem reads the four neighbor pixel values and computes their weighted sum.

```
open_system('ImageWarpHDL/ImageWarpHDLAlgorithm','force');
```
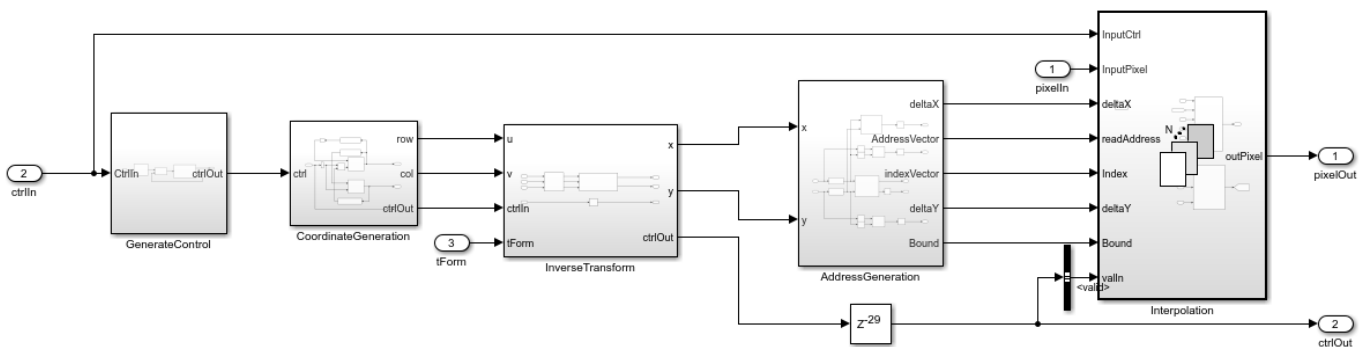

Inverse Transformation

The HDL implementation of the inverse geometric transformation multiplies the coordinates *[u v 1]* by the transformation matrix. The `Transformation` subsystem implements the matrix multiplication with Product blocks, which multiply the integer coordinates of the output image by each element of the transformation matrix. For this operation, the `Transformation` subsystem splits the transformation matrix into individual elements by using a Demux block. The `HomogeneousToCartesian` subsystem converts the generated homogeneous coordinates, *[x y z]* back to the cartesian format *[x y]* for further processing. The `HomogeneousToCartesian` subsystem uses a Reciprocal block configured to use the `ShiftAdd` architecture, and the Product blocks that compute *x* and *y* use the `ShiftAdd` architecture for better hardware clock speed. To see these parameters, right-click the block and select **HDL Code > HDL Block Properties**.

```
open_system('ImageWarpHDL/ImageWarpHDLAlgorithm/InverseTransform','force');
```
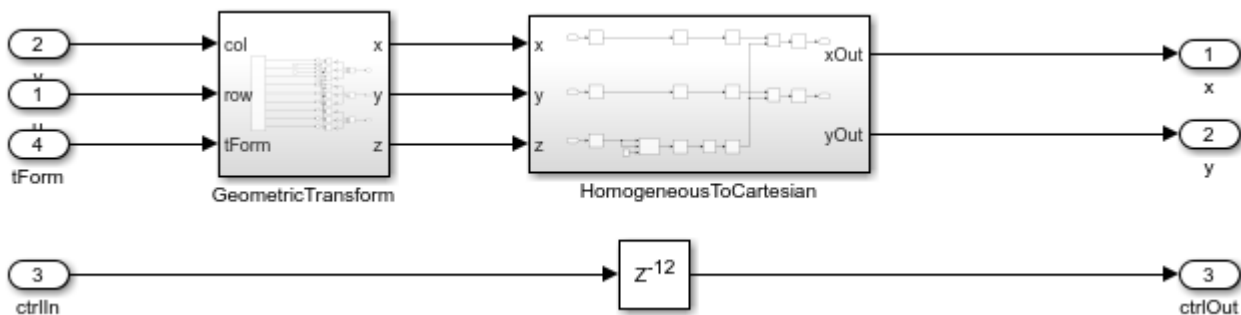
Address Generation

The `AddressGeneration` subsystem calculates the displacement of each pixel from its neighboring pixels by using the mapped coordinate *(x,y)* of the input raw image. The subsystem also rounds the coordinates to the nearest integer toward negative infinity.

```
open_system('ImageWarpHDL/ImageWarpHDLAlgorithm/AddressGeneration','force');
```
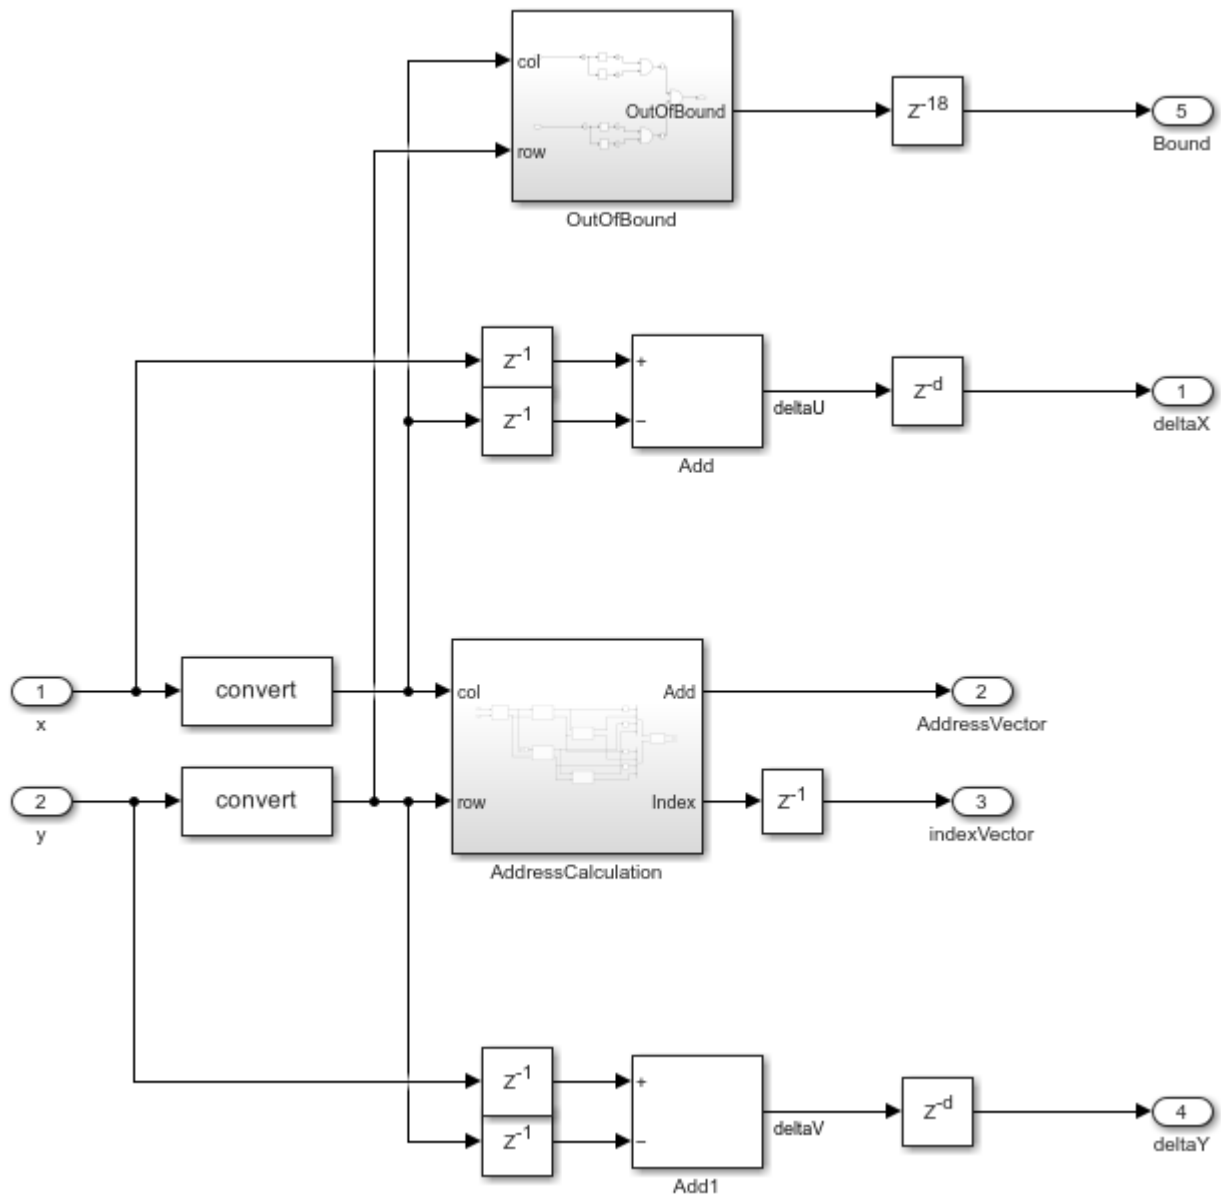

The `AddressCalculation` subsystem checks the coordinates against the bounds of the input images. If any coordinate is outside the image dimensions, the subsystem sets that coordinate to the boundary value. Next, the subsystem calculates the index of the address for each of the four neighborhood pixels in the `CacheMemory` subsystem. The index represents the column of the cache.

The subsystem finds the index for each address by using the even and odd nature of the incoming column and row coordinates, as determined by the Extract Bits block.

```
% ==========================
% |Row   || Col   || Index ||
% ==========================
% |Odd   || Odd   ||   1   ||
% |Even  || Odd   ||   2   ||
% |Odd   || Even  ||   3   ||
% |Even  || Even  ||   4   ||
% ==========================
```

This equation specifies the address of the neighborhood pixels.

$$addr = (\frac{colSize}{2} * R) + C$$

R is the row coordinate and C is the column coordinate. When `row` is even, then $nR = \frac{row}{2} - 1$. When `row` is odd, then $R = \frac{row-1}{2}$. When `col` is even, then $C = \frac{col}{2}$. When `col` is odd, then $C = \frac{col+1}{2}$.

The IndexChangeForMemoryAccess MATLAB Function block in the `AddressCalculation` subsystem rearranges the addresses in increasing order of their indices. This operation ensures the correct fetching of data from the CacheMemory block. This subsystem passes the addresses to the `CacheMemory` subsystem, and passes `Index`, ΔX, and ΔY to the `Interpolation` subsystem.

The `OutOfBound` subsystem checks whether the *(x,y)* coordinates are out of bounds (that is, if any coordinate is outside the image dimensions). If the coordinate is out of bounds, the subsystem sets the corresponding output pixel to an intensity value of `0`.

Finally, a Vector Concatenate block creates vectors of the addresses and indices.

Interpolation

The `Interpolation` subsystem is a For Each block, which replicates its operation depending on the dimensions of the input pixel. For example, if the input is an RGB image, then the input pixel dimensions are 1-by-3, and the model includes three instances of this operation. Because the model uses a For Each block, it supports RGB or grayscale input. The operation inside the `Interpolation` subsystem comprises two subsystems: `BilinearInterpolation` and `CacheMemory`.

```
open_system('ImageWarpHDL/ImageWarpHDLAlgorithm/Interpolation','force');
```

Cache Memory

The `CacheMemory` subsystem contains a Simple Dual Port RAM block. The subsystem buffers the input pixels to form [Line 1 Pixel 1 | Line 2 Pixel 1 | Line 1 Pixel 2 | Line 2 Pixel 2] in the RAM. By using this configuration, the algorithm can read all four neighboring pixels in one cycle. The example calculates the required size of the cache memory from the *offset* output of the `ComputeImageWarpCacheOffset` function. The offset is the sum of the *maximum deviation* and the *first row map*. The *first row map* is the maximum value of the input image row coordinate that corresponds to the first row of the output undistorted image. The *maximum deviation* is the greatest difference between the maximum and minimum row coordinates for each row of the input image row map.

The `WriteControl` subsystem forms vectors of incoming pixels, write enables, and write addresses. The `AddressGeneration` subsystem provides a vector of read addresses. The vector of pixels from the RAM is the input to the `BilinearInterpolation` subsystem.

```
open_system('ImageWarpHDL/ImageWarpHDLAlgorithm/Interpolation/CacheMemory','force');
```

Bilinear Interpolation

The `BilinearInterpolation` subsystem rearranges the vector of read pixels from the cache to their original indices. Then, the `BilinearInterpolationEquation` subsystem calculates a weighted sum of the neighborhood pixels. The result of the interpolation is the value of the output warped pixel.

```
open_system('ImageWarpHDL/ImageWarpHDLAlgorithm/Interpolation/BilinearInterpolation','force');
```
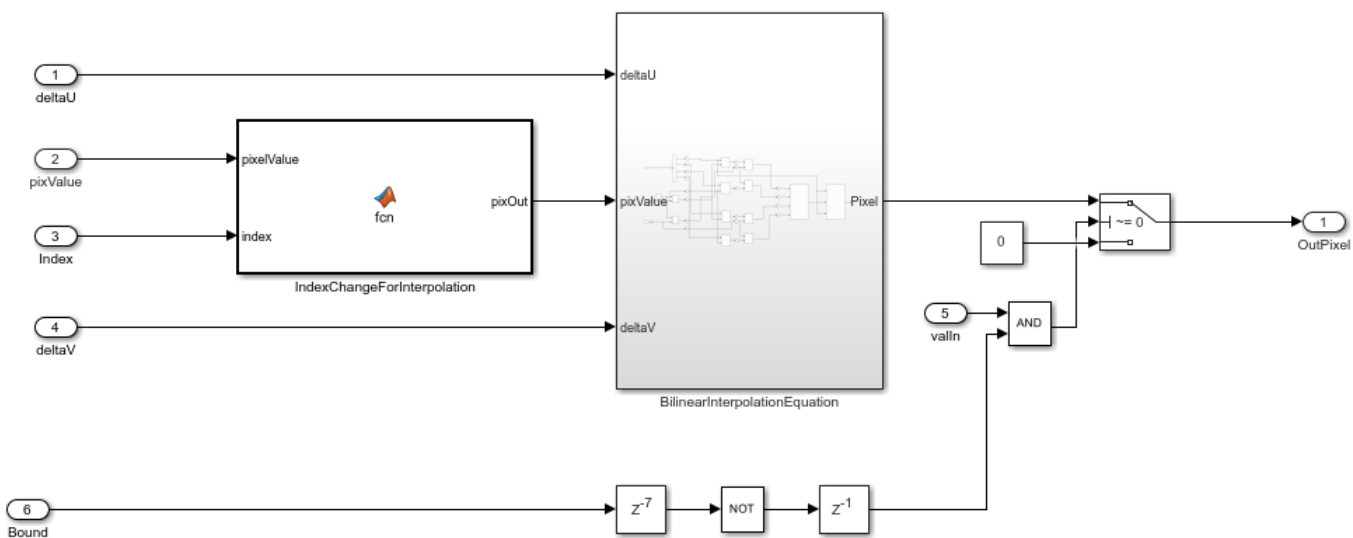

In the equation and the diagram, *(u,v)* is the coordinate of the input pixel generated by the inverse tranformation stage. I_1, I_2, I_3, and I_4 are the four neighboring pixels, and ΔU and ΔV are the displacements of the target pixel from its neighboring pixels. This stage of the algorithm computes the weighted average of the four neighboring pixels by using this equation.

$$outputPixel = I_1(1 - \Delta U)(1 - \Delta V) + I_2(\Delta U)(1 - \Delta V) + I_3(1 - \Delta U)(\Delta V) + I_4(\Delta U)(\Delta V)$$
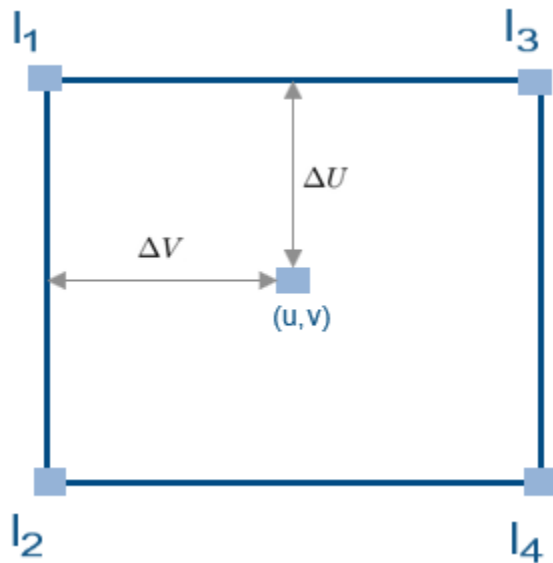
Simulation and Results

This example uses a 480p RGB input image. The input pixels use the `uint8` data type for either grayscale and RGB input images.

This implementation uses on chip BRAM memory rather than external DDR memory. The amount of BRAM required for the computation of output pixel intensities is directly proportional to the number of input lines required to be buffered in the cache. This bar graph shows the number of lines required in the cache for different angles of rotation of the output image. For this graph, the scaling factor is 1.1, and the translation in the *x*- and *y*-directions is 0.6 and 0.3, respectively.
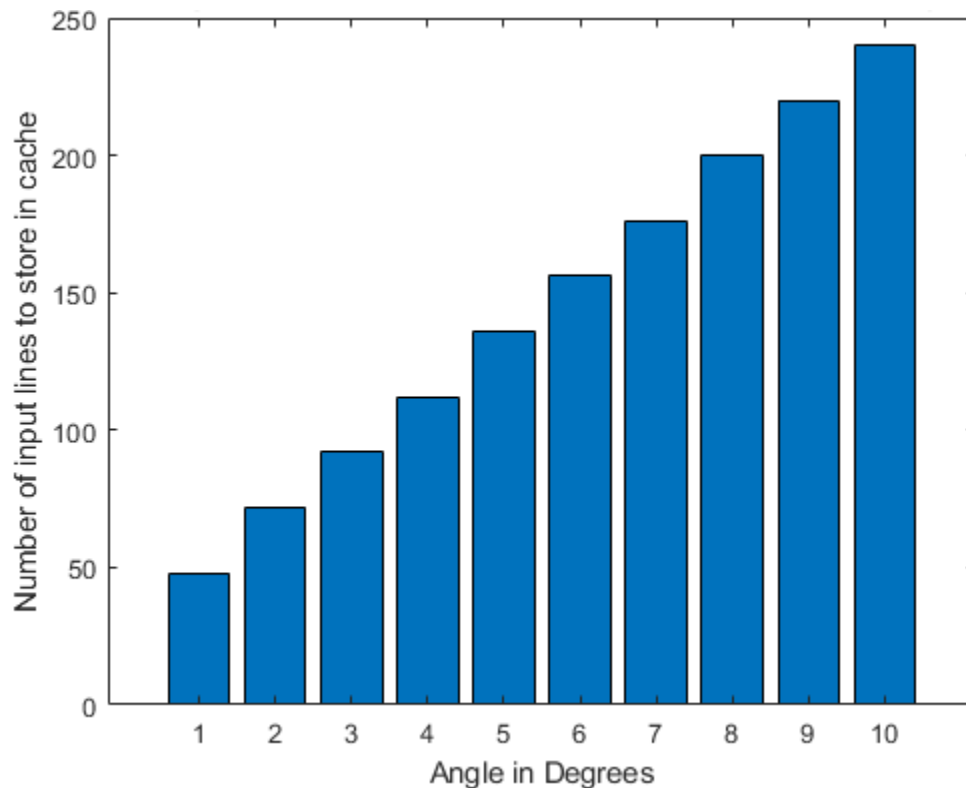
This figure shows the input image and the corresponding output image rotated by an angle of four degrees, scaled by a factor of 1.1, and translated by 0.4 and 0.8 in the *x*- and *y*-directions, respectively. The results of the ImageWarpHDL model match the output of the `imwarp` function in MATLAB.
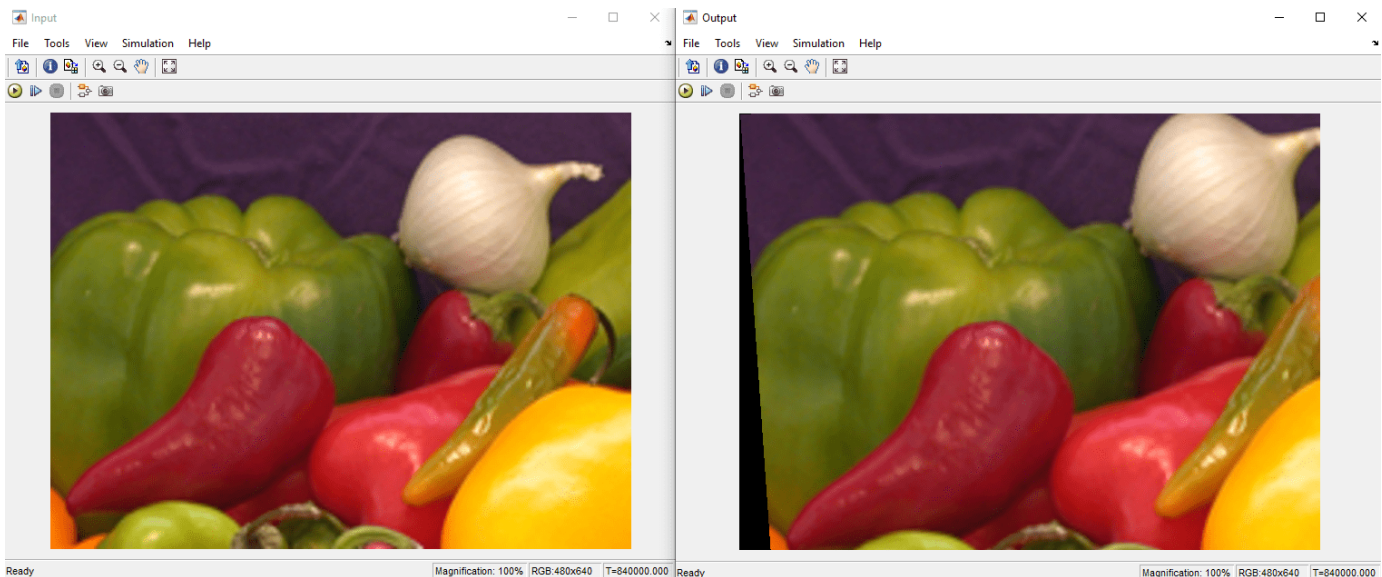
To check and generate the HDL code referenced in this example, you must have an HDL Coder™ license.

To generate the HDL code, use this command.

```
makehdl('ImageWarpHDL/ImageWarpHDLAlgorithm')
```

To generate the test bench, use this command.

```
makehdltb('ImageWarpHDL/ImageWarpHDLAlgorithm')
```

This design was synthesized using Xilinx® Vivado® for the Xilinx® Zynq®-7000 SoC ZC706 development kit and met a timing requirement of over 200MHz. The table shows the resource utilization for the HDL subsystem.

```
% ================================================================
% |Model Name            ||           ImageWarpHDL          ||
% ================================================================
% |Input Image Resolution ||           480 x 640            ||
% |Slice LUTs             ||           7325                 ||
% |Slice Registers        ||           7431                 ||
% |BRAM                   ||           97                   ||
% |Total DSP Blocks       ||           82                   ||
% ================================================================
```

See Also

More About

- "Stereo Image Rectification" on page 2-126
- "Image Undistortion" on page 2-136

Low Light Enhancement

This example shows how to enhance low-light images using an algorithm suitable for FPGAs.

Low-light enhancement (LLE) is a pre-processing step for applications in autonomous driving, scientific data capture, and general visual enhancement. Images captured in low-light and uneven brightness conditions have low dynamic range with high noise levels. These qualities can lead to degradation of the overall performance of computer vision algorithms that process such images. This algorithm improves the visibility of the underlying features in an image.

The example model includes a floating-point frame-based algorithm as a reference, a simplified implementation that reduces division operations, and a streaming fixed-point implementation of the simplified algorithm that is suitable for hardware.

LLE Algorithm

This example performs LLE by inverting an input image and then applying a de-haze algorithm on the inverted image. After inverting the low-light image, the pixels representing non-sky region have low intensities in at least one color channel. This characteristic is similar to an image captured in hazy weather conditions [1]. The intensity of these dark pixels is mainly due to scattering, or airlight, so they provide an accurate estimation of the haze effects. To improve the dark channel in an inverted low-light image, the algorithm modifies the airlight image based on the ambient light conditions. The airlight image is modified using the dark channel estimation and then refined with a smoothing filter. To avoid noise from over-enhancement, the example applies non-linear correction to better estimate the airlight map. Although this example differs in its approach, for a brief overview of low-light image enhancement, see "Low-Light Image Enhancement" (Image Processing Toolbox).

The LLE algorithm takes a 3-channel low-light RGB image as input. This figure shows the block diagram of the LLE Algorithm.
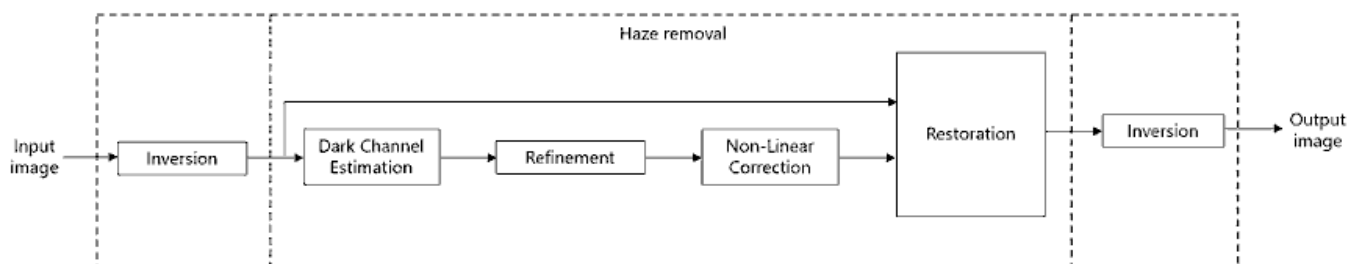
The algorithm consists of six stages.

1. *Scaling and Inversion*: The input image $I^c(x, y), c \in [r, g, b]$ is converted to range [0,1] by dividing by 255 and then inverting pixel-wise.

$$I^c_{scal}(x, y) = \frac{I^c(x, y)}{255}$$

$$I^c_{inv}(x, y) = 1 - I^c_{scal}(x, y)$$

2. *Dark Channel Estimation*: The dark channel is estimated by finding the pixel-wise minimum across all three channels of the inverted image [2]. The minimum value is multiplied by a haze factor, z, that represents the amount of haze to remove. The value of z is between 0 and 1. A higher value means more haze will be removed from the image.

$$I_{air}(x,y) = z \times \min_{c \, \epsilon \, [r,g,b]} I_{inv}^c(x,y)$$

3. *Refinement*: The airlight image from the previous stage is refined by iterative smoothing. This smoothing strengthens the details of the image after enhancement. This stage consists of five filter iterations with a 3-by-3 kernel for each stage. The refined image is stored in $I_{refined}(x,y)$. These equations derive the filter coefficients, h, used for smoothing.

$$I_{refined(n+1)}(x,y) = I_{refined(n)}(x,y) * h, \ n = [0,1,2,3,4] \ \& \ I_{refined(0)} = I_{air}$$

$$where \ h = \frac{1}{16} \begin{bmatrix} 1 & 2 & 1 \\ 2 & 4 & 2 \\ 1 & 2 & 1 \end{bmatrix}$$

$$Let \ I_{refined(5)}(x,y) = I_{refined}(x,y)$$

4. *Non-Linear Correction*: To reduce over-enhancement, the refined image is corrected using a non-linear correction equation shown below. The constant, m, represents the mid-line of changing the dark regions of the airlight map from dark to bright values. The example uses an empirically-derived value of $m = 0.6$.

$$I_{nlc}(x,y) = \frac{[I_{refined}(x,y)]^4}{[I_{refined}(x,y)]^4 + m^4}$$

5. *Restoration*: Restoration is performed pixel-wise across the three channels of the inverted and corrected image, I_{nlc}, as shown:

$$I_{restore}^c(x,y) = \frac{I_{scal}^c(x,y) - I_{nlc}(x,y)}{1 - I_{nlc}(x,y)}$$

6. *Inversion*: To obtain the final enhanced image, this stage inverts the output of the restoration stage, and scales to the range [0,255].

$$I_{enhanced}^c(x,y) = 255 \times (1 - I_{restore}^c)$$

LLE Algorithm Simplification

The scaling, non-linear correction, and restoration steps involve a divide operation which is not efficient to implement in hardware. To reduce the computation involved, the equations in the algorithm are simplified by substituting the result of one stage into the next stage. This substitution results in a single constant multiplication factor rather than several divides.

Dark channel estimation without scaling and inversion is given by

$$I_{air}(x,y) = \frac{z}{255} I_{air}'(x,y) \text{ where } I_{air}'(x,y) = 255 - \min_{c \, \epsilon \, [r,g,b]} I^c(x,y)$$

The result of the iterative refinement operation on I_{air} is

$$I_{refined}(x,y) = \frac{z}{255} I'_{refined(5)}(x,y)$$

where

$$I'_{refined(n+1)}(x,y) = I'_{refined(n)}(x,y) * h, \ n = [0,1,2,3,4] \ \& \ I'_{refined(0)}(x,y) = I'_{air}(x,y)$$

Substituting $I_{refined}$ into the non-linear correction equation gives

$$I_{nlc}(x,y) = \frac{z^4[I'_{refined}(x,y)]^4}{z^4[I'_{refined}(x,y)]^4 + (255 \times m)^4}$$

Substituting I_{nlc} into the restoration equation gives

$$I^c_{restore}(x,y) = 1 - \frac{I^c(x,y)}{255} - \frac{I^c(x,y)}{255} \frac{z^4}{(255 \times m)^4}[I'_{refined}(x,y)]^4$$

Subtracting $I^c_{restore}$ from 1 and multiplying by 255 gives

$$I^c_{enhanced}(x,y) = I^c(x,y) \times \left(1 + \left[\frac{z}{255 \times m} I'_{refined}(x,y)\right]^4\right)$$

With the intensity midpoint, m, set to 0.6 and the haze factor, z, set to 0.9, the simplified equation is

$$I^c_{enhanced}(x,y) = I^c(x,y) \times \left(1 + \left[\frac{1}{170} I'_{refined}(x,y)\right]^4\right)$$

In the equation above, the factor multiplied with $I^c(x,y)$ can be called the Enhancement Factor. The constant $\frac{1}{170}$ can be implemented as a constant multiplication rather than a divide. Therefore, the HDL implementation of this equation does not require a division block.

HDL Implementation

The simplified equation is implemented for HDL code generation by converting to a streaming video interface and using fixed-point data types. The serial interface mimics a real video system and is efficient for hardware designs because less memory is required to store pixel data for computation. The serial interface also allows the design to operate independently of image size and format, and makes it more resilient to video timing errors. Fixed-point data types use fewer resources and give better performance on FPGA than floating-point types.

```
open_system('LLEExample');
```
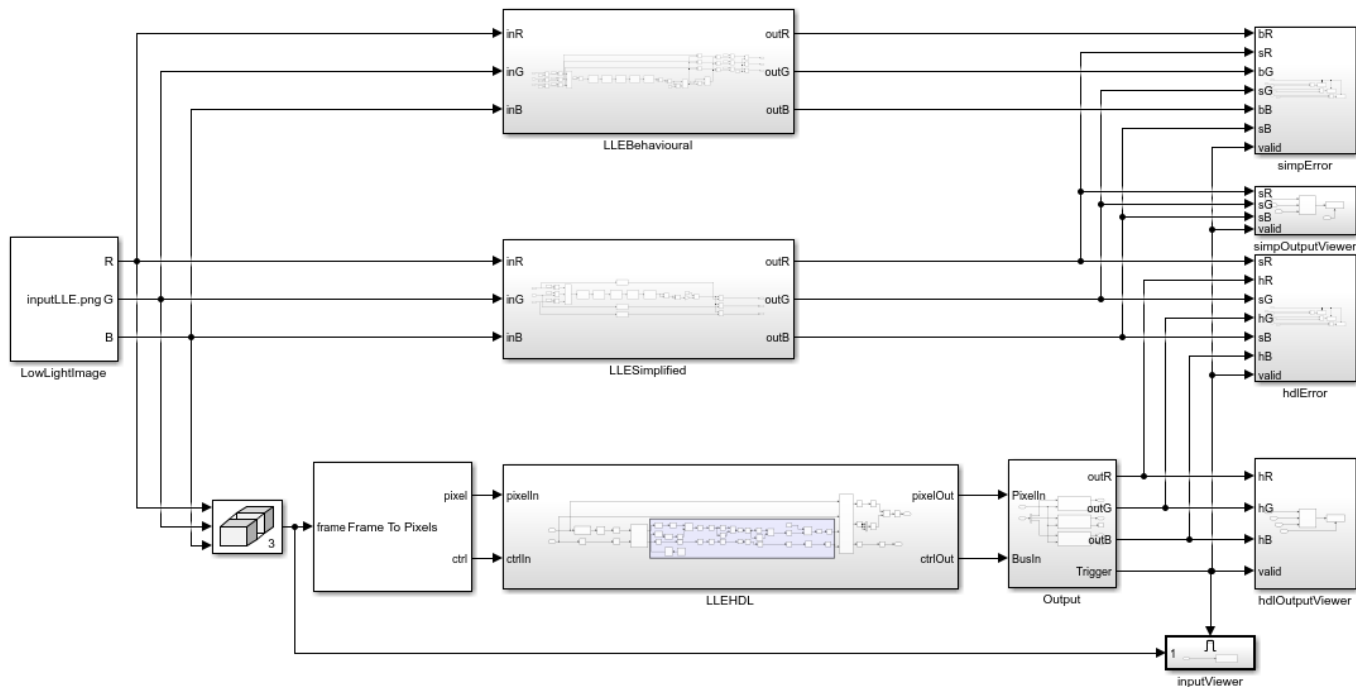
Copyright 2018 The MathWorks, Inc.

The location of the input image is specified in the *LowLightImage* block. The *LLEBehavioural* subsystem computes the enhanced image using the raw equations as described in the LLE Algorithm section. The *LLESimplified* subsystem computes the enhanced image using the simplified equations. The *simpOutputViewer* shows the output of the *LLESimplified* subsystem.

The *LLEHDL* subsystem implements the simplified equation using a steaming pixel format and fixed-point blocks from Vision HDL Toolbox™. The *Input* subsystem converts the input frames to a pixel stream of `uint8` values and a `pixelcontrol` bus using the Frame To Pixel block. The *Output* subsystem converts the output pixel stream back to image frames for each channel using the Pixel To Frame block. The resulting frames are compared with result of the *LLESimplified* subsystem. The *hdlOutputViewer* subsystem and *inputViewer* subsystem show the enhanced output image and the low-light input image, respectively.

```
open_system('LLEExample/LLEHDL');
```
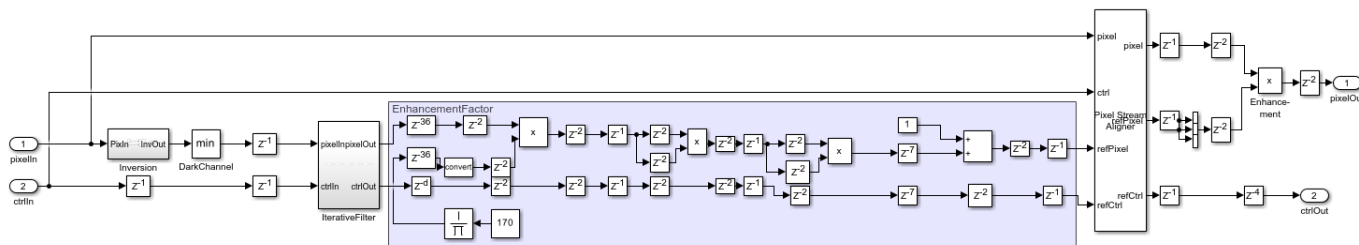

The *LLEHDL* subsystem inverts the input `uint8` pixel stream by subtracting each pixel from 255. Then the *DarkChannel* subsystem calculates the dark channel intensity minimum across all three channels. The *IterativeFilter* subsystem smooths the airlight image using sequential Image Filter blocks. The bit growth of each filter stage is maintained to preserve the precision. The Enhancement

Factor is calculated in EnhancementFactor area. The constant $\frac{1}{170}$ is implemented using Constant and Reciprocal blocks. The Pixel Stream Aligner block aligns the input pixel stream with the pipelined, modified stream. The aligned input stream is then multiplied by the modified pixel stream.

Simulation and Results

The input to the model is provided in the *LowLightImage* (Image From File) block. This example uses a 720-by-576 pixel input image with RGB channels. Both the input pixels and the enhanced output pixels use `uint8` data type. The necessary variables for the example are initialized in PostLoadFcn callback.

The *LLEBehavioural* subsystem uses floating-point Simulink blocks to prototype the equations mentioned in the LLE Algorithm section. The *LLESimplified* subsystem implements the simplified equation in floating-point blocks, with no divide operation. The *LLEHDL* subsystem implements the simplified equation using fixed-point blocks and streaming video interface. The figure shows the input image and the enhanced output images obtained from the *LLESimplified* subsystem and the *LLEHDL* subsystem.
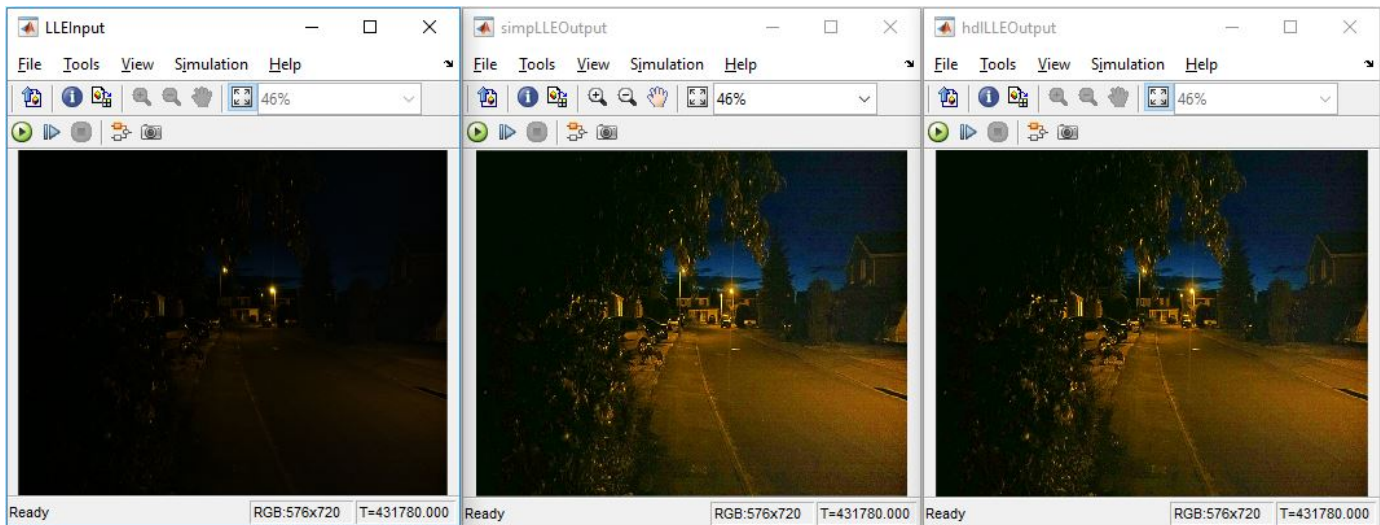
The accuracy of the result can be calculated using the percentage of error pixels. To compute the percentage of error pixels in the output image, the difference between the pixel value of the reference output image and the *LLEHDL* output image should not be greater than one, for each channel. The percent of pixel values that differ by more than 1 is computed for the three channels. The *simpError* subsystem compares the result of the *LLEBehavioural* subsystem with the result of the *LLESimplified* subsystem. The *hdlError* subsystem compares the result of the *LLEHDL* subsystem with the result of the *LLESimplified* subsystem. The error pixel count is displayed for each channel. The table shows the percentage of error pixels calculated by both comparisons.

	% error pixels		
Model Name / Channel	R	G	B
LLEBehavioural vs LLESimplified	0	0	0
LLESimplified vs LLEHDL	0.1649	0.1671	0.0477

References

[1] X. Dong, G. Wang, Y. Pang, W. Li, and J. Wen, "Fast efficient algorithm for enhancement of low lighting video" IEEE International Conference on Multimedia and Expo, 2011.

Contrast Limited Adaptive Histogram Equalization

This example shows how to implement a contrast-limited adaptive histogram equalization (CLAHE) algorithm using Simulink® blocks. The example model is FPGA-hardware compatible.

The example uses the `adapthisteq` function from the Image Processing Toolbox™ as reference to verify the design.

Introduction

Adaptive histogram equalization (AHE) is an image pre-processing technique used to improve contrast in images. It computes several histograms, each corresponding to a distinct section of the image, and uses them to redistribute the luminance values of the image. It is therefore suitable for improving the local contrast and enhancing the definitions of edges in each region of an image. However, AHE has a tendency to overamplify noise in relatively homogeneous regions of an image. A variant of adaptive histogram equalization called contrast-limited adaptive histogram equalization (CLAHE) prevents this effect by limiting the amplification.

CLAHE Algorithm
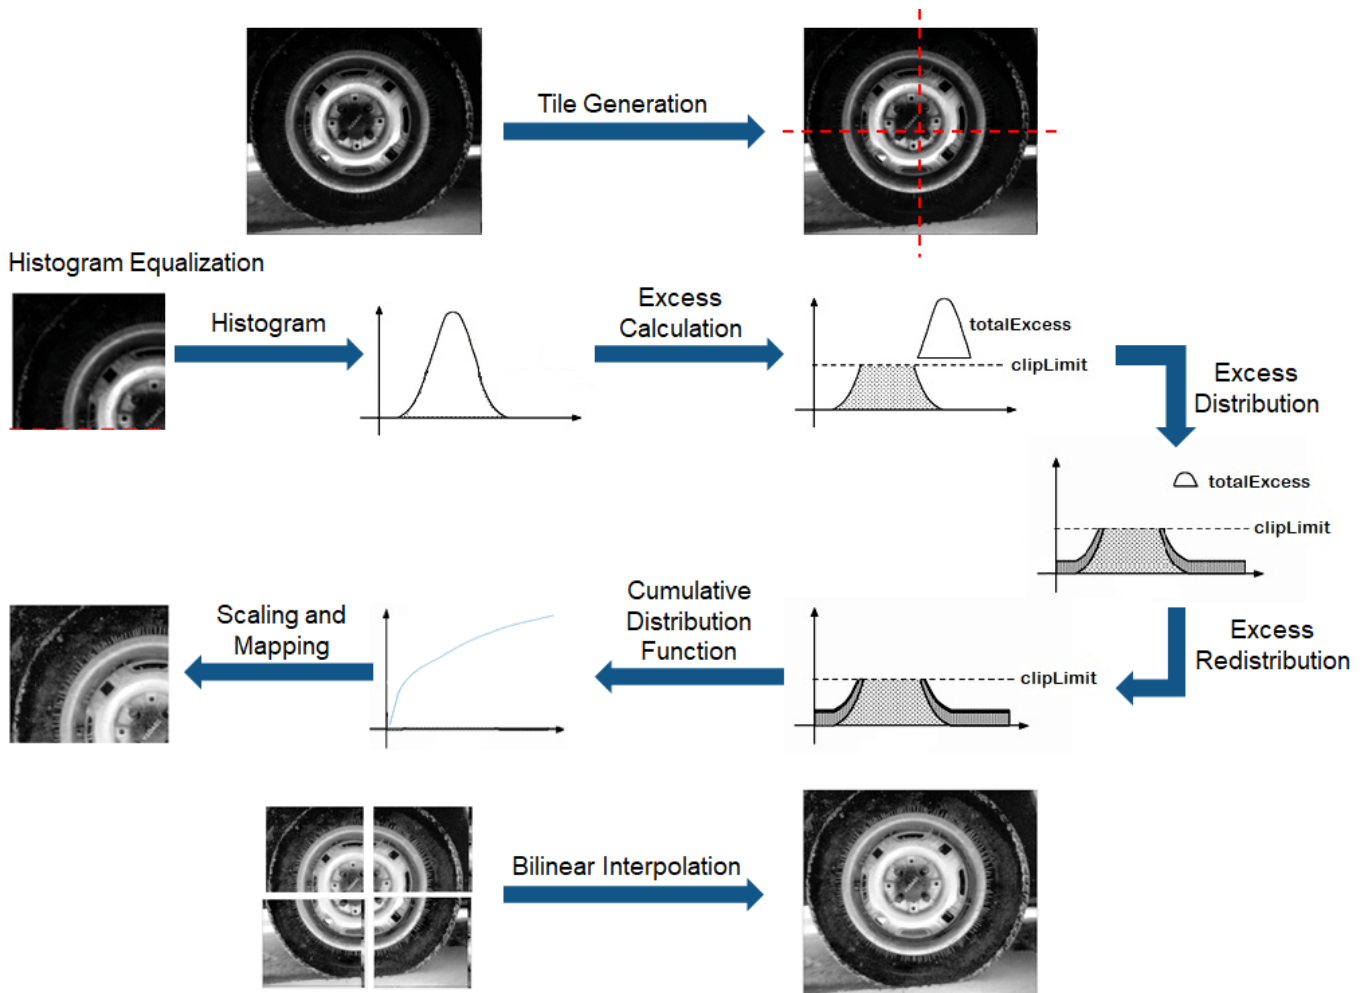
The CLAHE algorithm has three major parts: tile generation, histogram equalization, and bilinear interpolation. The input image is first divided into sections. Each section is called a tile. The input image shown in the figure is divided into four tiles. Histogram equalization is then performed on each tile using a pre-defined clip limit. Histogram equalization consists of five steps: histogram computation, excess calculation, excess distribution, excess redistribution, and scaling and mapping using a cumulative distribution function (CDF). The histogram is computed as a set of bins for each tile. Histogram bin values higher than the clip limit are accumulated and distributed into other bins. CDF is then calculated for the histogram values. CDF values of each tile are scaled and mapped using the input image pixel values. The resulting tiles are stitched together using bilinear interpolation, to generate an output image with improved contrast.

HDL Implementation
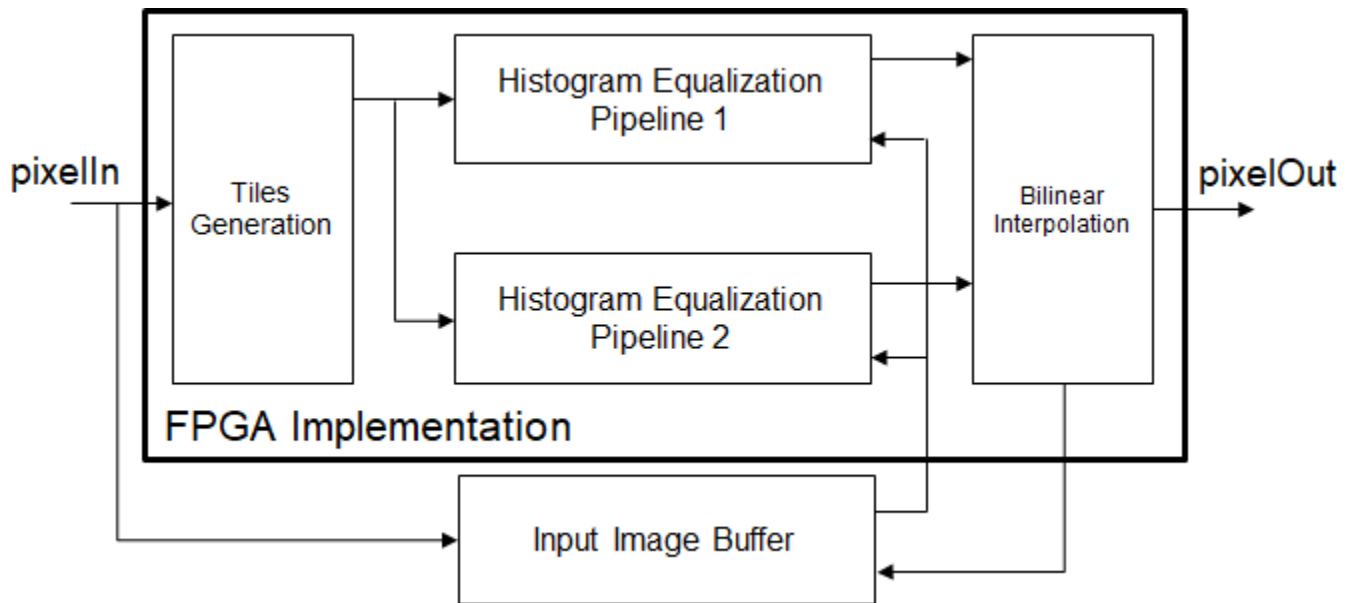
This figure shows the block diagram of the HDL implementation of the CLAHE algorithm. It consists of a tile generation block, a histogram equalization pipeline block, a bilinear interpolation block, and an input image buffer block. Tiles are generated by modifying the `pixelcontrol` bus of the pixel stream for the desired tile size. The pixel stream and the modified `pixelcontrol` bus are fed to the histogram equalization pipeline. Two histogram equalization pipelines are required to keep pace with the input data. They operate in ping-pong manner. Each pipeline contains histogram equalization modules equal to the number of tiles in the horizontal direction. The histogram equalization modules work in parallel to compute histogram equalization for each tile. The last stage in the histogram equalization module, scaling and mapping, needs the original input image data. This data is stored in an input image buffer block. The bilinear interpolation block generates addresses to read the input image pixel values from the memory. The input image pixel values from the image buffer block are given to the histogram equalization modules for mapping. Mapped values obtained from histogram equalization are scaled and used in the bilinear interpolation computation to reduce boundary artifacts.

```
modelname = 'CLAHEExample';
open_system(modelname,'force');
set_param(modelname,'SampleTimeColors','off');
set_param(modelname,'Open','on');
```

```
set_param(modelname,'SimulationCommand','Update');
set(allchild(0),'Visible','off');
```
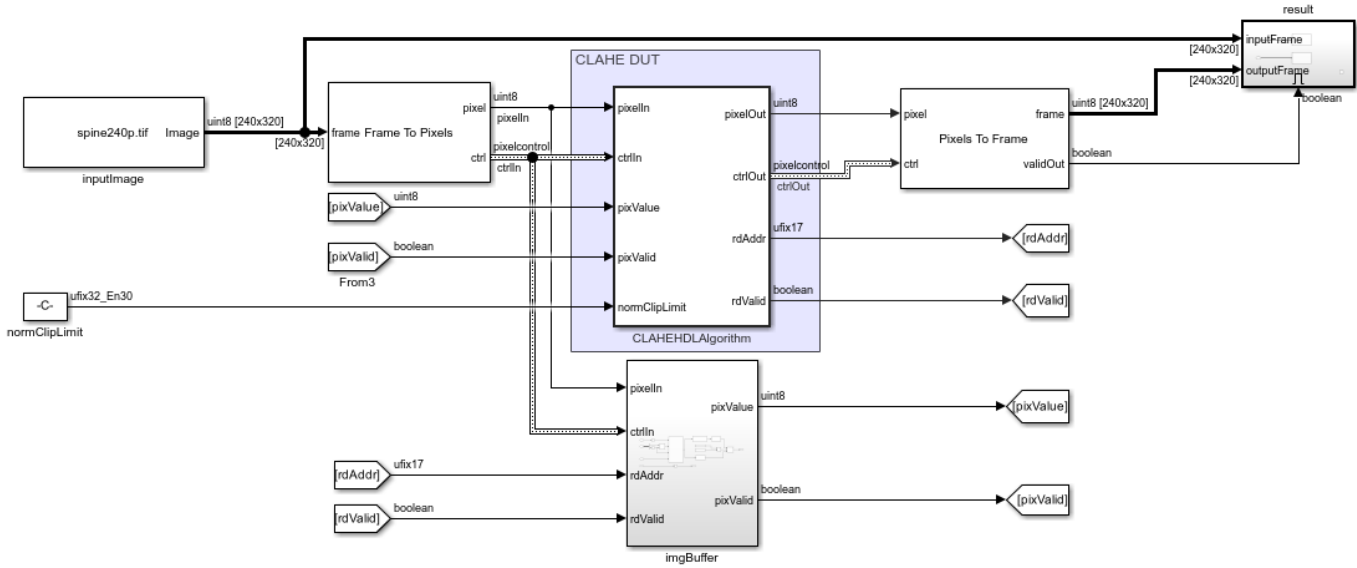

Copyright 2023 The MathWorks, Inc.

The figure shows the top level view of the `CLAHEExample` model. The input image path is specified in the inputImage block. The input image frame is converted to a pixel stream and `pixelcontrol` bus using a Frame To Pixels block. The pixel stream is passed to the CLAHEHDLAlgorithm subsystem for contrast enhancement and is also stored in the imgBuffer subsystem. While processing, the CLAHEHDLAlgorithm subsystem generates the address to read image data from the imgBuffer subsystem. The pixel value read from the imgBuffer subsystem is passed to CLAHEHDLAlgorithm for adjustment. The adjusted pixel values are given to the Pixels To Frame block and converted to a frame using the control signals. The Result subsystem shows the input image and output image once all the pixels in the frame have been received by the Pixels To Frame block.

Tile Generation

```
system = 'CLAHEExample/CLAHEHDLAlgorithm/tileGeneration';
open_system(system,'force');
```

Tile generation subsystem is used to define the contextual regions for contrast enhancement by dividing the input image across horizontal and vertical directions
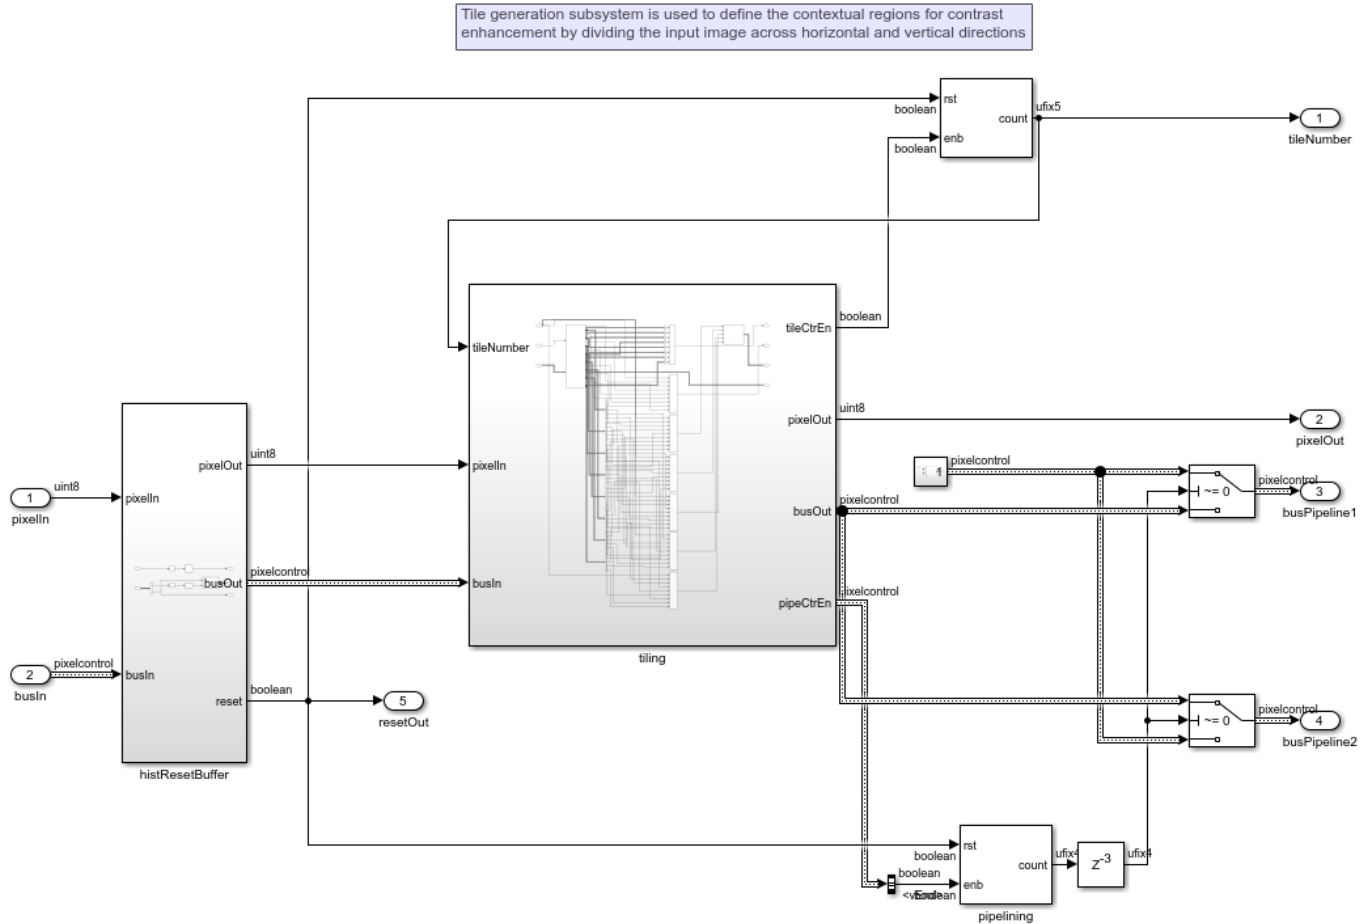
The figure shows the tile generation subsystem. This subsystem is used to divide the input image into a number of tiles in both the horizontal and vertical directions. By default, the model divides the input image into 8 tiles in each direction. Tiles are created by modifying the input `pixelcontrol` bus to select the pixels in each tile region. The size of a vertical(horizontal) tile is computed by dividing the number of rows(columns) in the input image by the number of tiles in the same direction. Inside the tiling subsystem, the ROI Selector block has vertical reuse enabled. This option enables parallel processing of the vertical tiles and the RoI Selector generates `pixel` streams and corresponding `pixelcontrol` buses for each of the horizontal tiles. The pixel stream to the histogram equalization pipeline is controlled by diverting each vertical tile to an alternate pipe. The tile size calculated in either must be an even integer. If the input image does not divide into an integer number of even-sized tiles, pad the input image symmetrically.

Histogram Equalization Pipeline

```
system = 'CLAHEExample/CLAHEHDLAlgorithm/histoEqPipeline/';
subsystem = [system 'histPipe1'];
open_system(subsystem,'force');
```
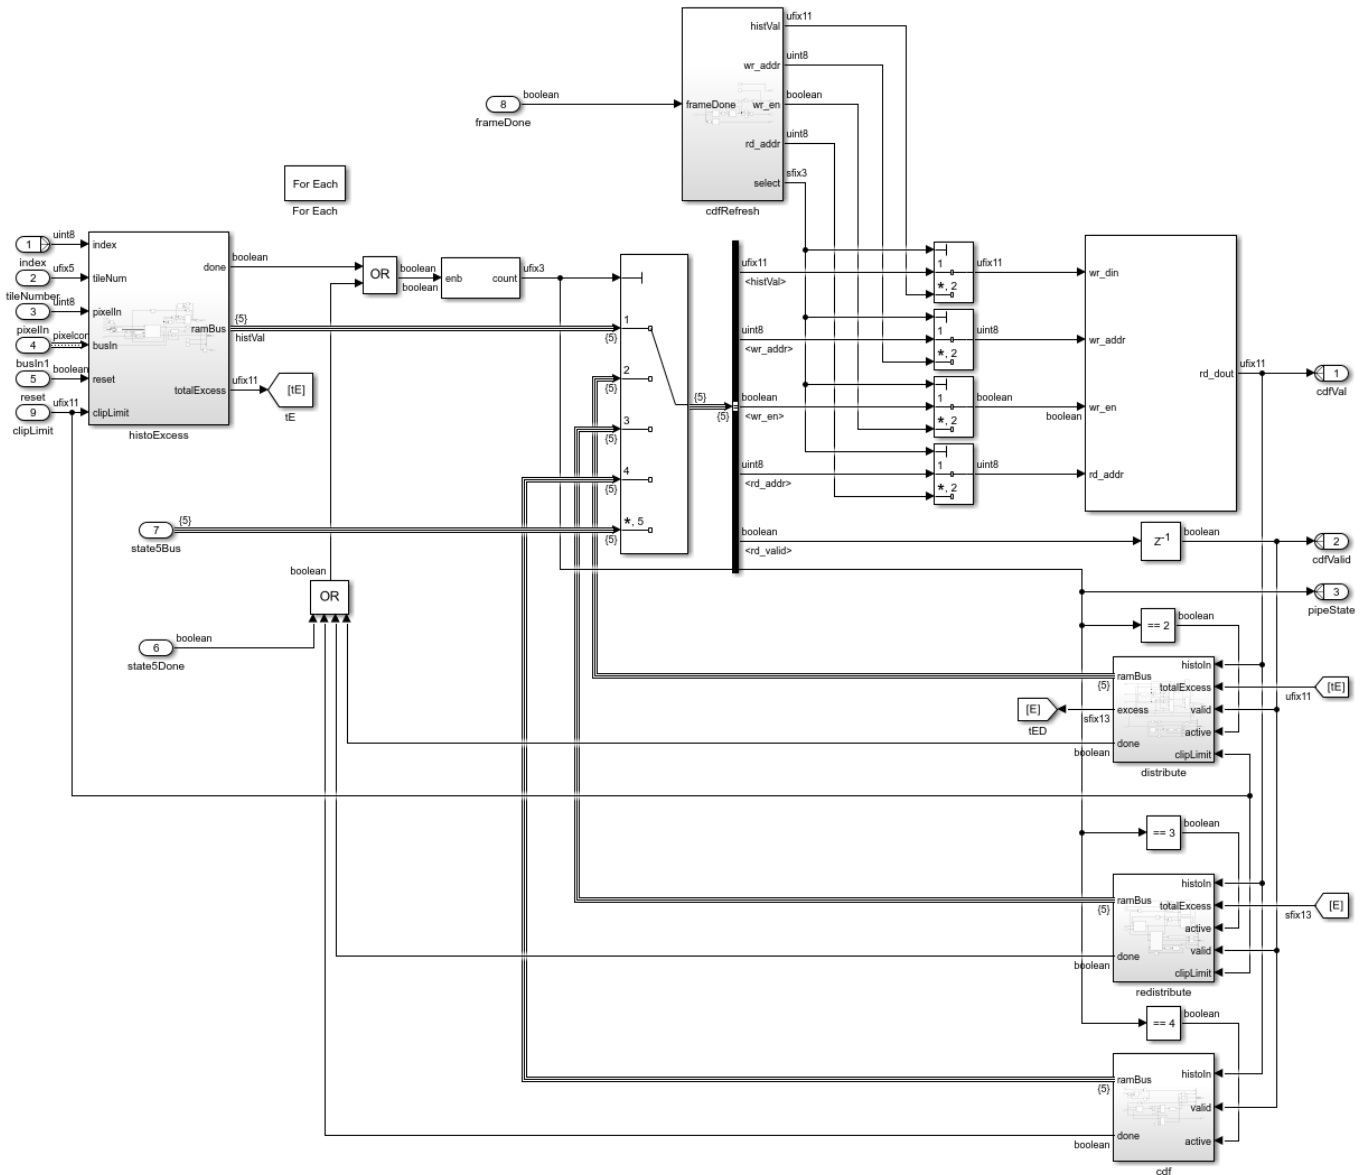
Two histogram equalization pipelines are used to keep pace with the streaming input pixels. Each histogram equalization pipeline consists of histogram equalization modules corresponding to each tile in the horizontal direction. These modules are implemented by using a For Each subsystem. Each histogram equalization module is divided into five stages: histogram calculation, total excess calculation, total excess distribution, excess redistribution, cumulative distribution function, and mapping.

The first module of the histogram pipeline, histoExcess subsystem, performs histogram calculation and total excess calculation for each tile. To compute the histogram, the Histogram block is used. When the histogram is complete the block generates a **readRdy** signal. The subsystem then reads the histogram values and determines excess value from each bin by using clip limit value. The clip limit is computed from the normalized clip limit value specified using these equations.

$$minClipLimit = ceil(numPixInTile/numBins);$$

$$clipLimit = minClipLimit + round(normClipLimit * (numPixInTile - minClipLimit));$$

The excess value from each bin is accumulated to form total excess value. The previously computed histogram values are not changed during total excess calculation and are stored in a Simple Dual Port RAM memory block. The necessary control signals for the RAM block (ramBus) are generated by the histoExcess subsystem. The total excess value calculated in the histoExcess subsystem is used by the Distribute subsystem.

The Distribute subsystem computes two variables: average bin increment and upper limit. These values are computed from the total excess value by using these equations:

$$avgBinIncr = totalExcess/numBins;$$

$$upperLimit = clipLimit - avgBinIncr;$$

The Distribute subsystem then reads the value of each histogram bin from the RAM block. It updates the value at every bin based on these three conditions:

1   If the histogram value of a bin is greater than the clip limit, it is replaced with the clip limit.
2   If the histogram value of a bin is between the clip limit and the upper limit, the histogram value is replaced with the clip limit. The total excess value is reduced by the number of added pixels equal to (`clipLimit - histVal`).
3   If the histogram value of a bin is less than the upper limit, the histogram value is increased by the average bin increment. The total excess value is reduced by the average bin increment.

The adjusted histogram value is stored at the same address. The remaining total excess value is passed to the Redistribute subsystem as excess value.

```
system = 'CLAHEExample/CLAHEHDLAlgorithm/histoEqPipeline/';
subsystem = [system 'histPipe1/redistribute'];
open_system(subsystem,'force');
```
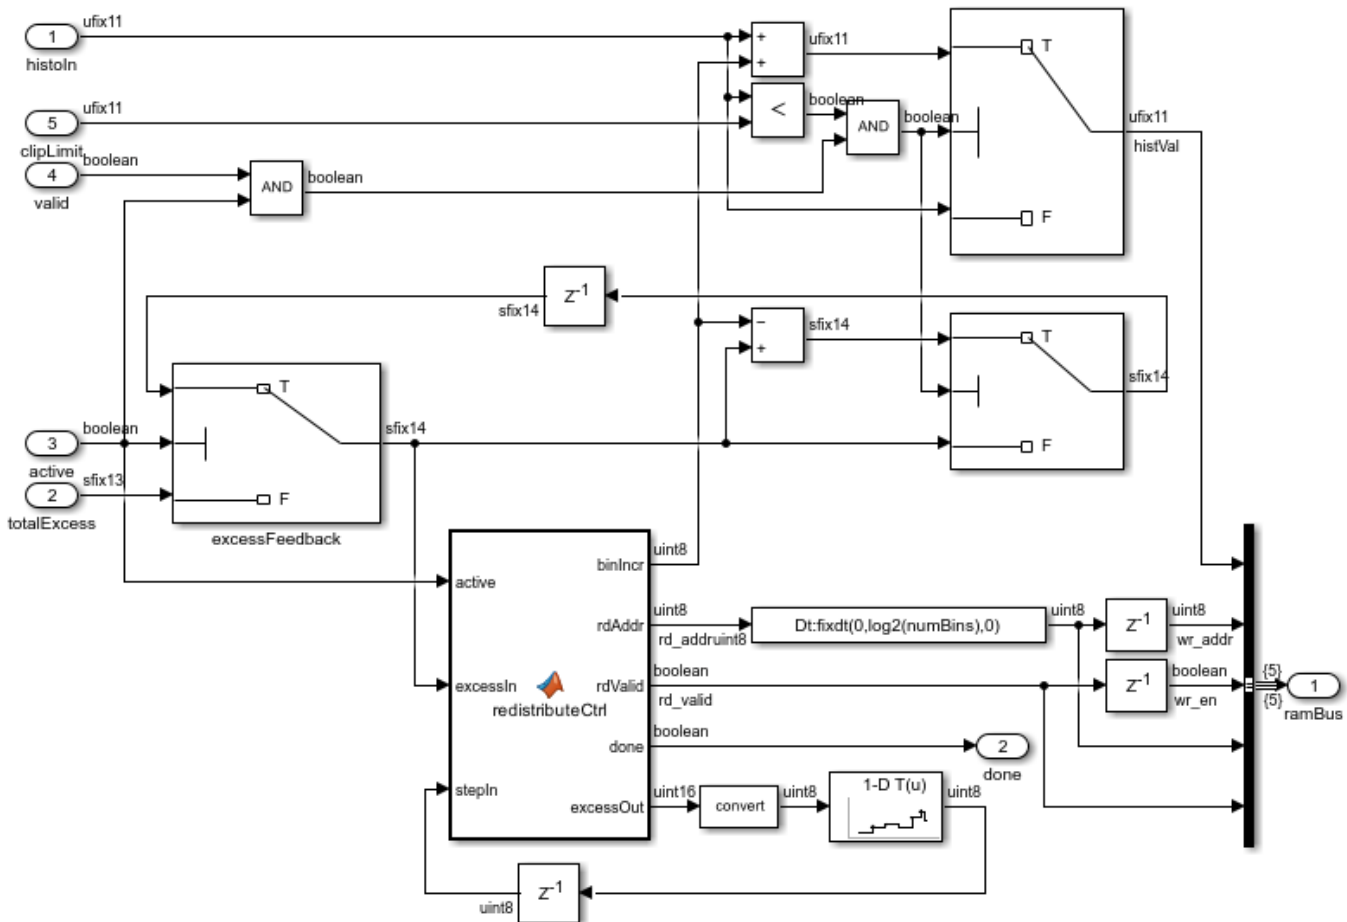
The Redistribute subsystem distributes spillover excess values to the histogram bins. It primarily uses two variables to distribute excess values: **binIncr** and **step**. **binIncr** specifies the value to be added to the histogram bins. **step** specifies the increment in the address counter used to fetch the histogram bin value. If the excess is greater than or equal to the number of bins, then **binIncr** is calculated by dividing the excess value by the number of bins, and **step** is set to 1. The divide is implemented by using a right-shift operation, since the number of bins is a power of 2.

If the excess is less than the number of bins, **binIncr** is set to 1 and **step** is calculated by dividing the number of bins by the excess value. The divide is computed by using a n-D Lookup Table (Simulink) block. The `redistributeCtrl` MATLAB Function generates the address for the RAM block by using the **step** value computed. When the address reaches the total number of bins, the **step** value is re-computed using the most recent excess value. Care is taken to not repeat the first bin as the start bin for redistribution. If the value of the histogram bin is less than the clip limit, it is increased by **binIncr**, and the same value is subtracted from the excess value. If the value of histogram bin is equal to the clip limit, no operation is performed and the value is written back to the same address. The MATLAB Function block repeats these bin adjustments until the excess value reaches 0.

The last stage of the histogram equalization pipeline is CDF calculation. The CDF subsystem computes the cumulative sum of the histogram bin values. The histogram values are read from the RAM block and added to the sum of the previous histogram bin values. It is then stored to the same address.

The five stages of the histogram equalization module can be considered as five states. The five states of histogram equalization module are sequential. Thus, a state counter is used to move from one state to another state. A counter value determines the state of the histogram equalization module. A Multiport Switch (Simulink) block is used with the state counter as the index value. The multi-port switch connects the ramBus from each state with the correct memory according to the index. The state counter is in state 1 in idle condition. When histoExcess finishes excess calculation it sets the **done** signal to 1 for one cycle, and the state counter moves to state 2. Similarly, the distribute subsystem, redistribute subsystem, and cdf subsystem generate done flags when their processing completes. These done flags increment the state counter to state 5, where it uses input image pixel values from the input image buffer block as addresses to read CDF values from the RAM. Before being used as address, the input image pixel values are scaled according to the number of histogram bins. When the number of histogram bins are less than the number of input image intensity levels, the latter values are mapped to the same range as CDF values. The address counter that reads the input image values is driven by the bilinear interpolation subsystem. The state counter is incremented by the bilinear interpolation subsystem when mapping for the respective pipeline is complete.

Bilinear Interpolation
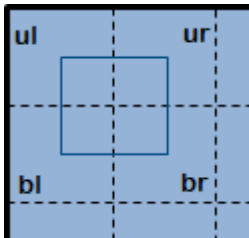
Bilinear interpolation is used to smooth edges when the tiles are stitched together. The figure shows how four tiles are used to compute pixel values in the output image. The each tile is divided into four parts. One part from each of the four tiles are grouped together to compute bilinear interpolation for that section of the image.

Interpolation uses this equation:

```
claheI(imgTileIdx{1}, imgTileIdx{2}) = ...
    (rowRevW .* (colRevW .* double(grayxform(imgPixVals,ulMapTile)) + ...
              colW     .* double(grayxform(imgPixVals,urMapTile)))+ ...
     rowW     .* (colRevW .* double(grayxform(imgPixVals,blMapTile)) + ...
              colW     .* double(grayxform(imgPixVals,brMapTile))))...
    /normFactor;
```

$$grayxform(imgPixVals, mapTile) = round((2^{inputBitWidth} - 1) * mapTile(imgPixVals)/numPixInTile);$$

The bilinear interpolation equation uses the position of a pixel with respect to each tile and the intensity information at that position to compute a pixel value in the output image. The intensity information is obtained from the input image pixel values stored in the image buffer. For corner tiles, intensity values are replicated (mirrored). The intensity information at the respective position in each tile is extracted from the CDF function of the histogram equalization pipeline by using the input image pixel value at the same position. The `grayxform` function scales the values obtained from the CDF function. The result is then divided by the number of pixels in a tile, represented as *normFactor* in the equation.

```
system = 'CLAHEExample/CLAHEHDLAlgorithm/bilinearInterpolation';
open_system(system,'force');
```


The figure shows the HDL implementation of the bilinear interpolation subsystem. When the histogram equalization pipeline reaches state 5, the paramCalc subsystem starts computing the read address for the imgBuffer subsystem. The pixel value read from the buffered image is the address for the RAM in the histogram equalization pipeline. CDF values are fetched from the read address for all the tiles from both the histogram equalization pipelines simultaneously. The required CDF values are selected and passed to the equation subsystem using Selector Switch blocks and Switch blocks. The Switch block selects which pipeline contains upper/lower tiles and the Selector Switch blocks select data corresponding to left/right tiles. The control signals for the Selector Switch and Switch blocks are generated in the paramCalc subsystem by using a read counter. Thus, intensity values at a pixel position for each tile are obtained from the image buffer. The bilinear interpolation equation also requires the pixel position and the total number of pixels in the tile. These parameters are also generated in the paramCalc subsystem. The equation subsystem is pipelined to optimize performance in hardware. The result is returned as a pixel stream with a `pixelcontrol` bus.
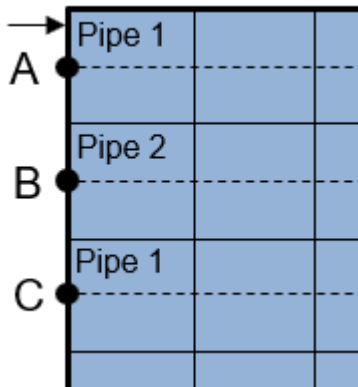
Bilinear interpolation of the output image is computed by traversing the rows from left to right. When all histogram equalization modules in the first pipeline have reached state 5, the paramCalc subsystem is enabled. The read addresses for the imgBuffer subsystem are computed until point A. Further computation of bilinear interpolation requires values from the histogram equalization modules of the second pipeline. When all histogram equalization modules in the second pipeline have reached state 5, the read address counter is again enabled and the bilinear interpolation output results are computed for pixel positions between point A and point B. Once the address counter reaches point B, results from first pipeline are no longer required. The `pipe1Done` signal is generated to change the state of the first histogram equalization pipeline modules back to state 1. Until this point, the tiles in the first pipeline are upper tiles and the tiles in the second pipeline are lower tiles. For the computation of values between point B and point C, the tiles in the second pipeline become the upper tiles and tiles in the first pipeline are now lower tiles. This operation continues until only the lowest tiles in the image remain. The output for these tiles is computed by replicating the values for the other pipeline. The output results are pushed into a FIFO in the outputStage subsystem and popped out such that the output valid signal is similar to that of the input pixel stream.

Model Parameters
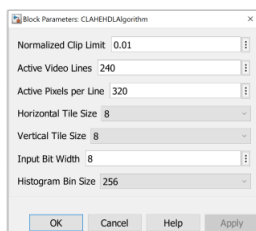
CLAHE uses a clip limit to prevent over-saturation of the image in homogeneous areas. These areas are characterized by a high peak in the histogram of an image tile due to many pixels falling in the same intensity range. For the model presented here, the clip limit is a user-defined normalized value. The default value is 0.01 (as shown in figure). The clip limit can be any value between 0 and 1 (inclusive).

The input image frame dimensions are specified by Active Video Lines and Active Pixels Per Line. The input image frame size is essential in setting the tile dimensions. Tiles define the number of rectangular contextual regions into which the image is divided. The horizontal and vertical tile size refer to the number of tiles in the relevant direction. Both these values must be at least 2 and the input image can only be divided into an integer number of even-sized tiles. The tile size mask

parameters are automatically populated with the valid options for each image dimension. The optimal number of tiles depends on the type of the input image, and it is best determined through experimentation.

The input bit width defines the number of bits per pixel in the input image. This helps to determine the maximum intensity value the input image can represent. The number of histogram bins used to build the contrast enhancing transformation can be varied from 32 to 4096. If the image dimensions or tile sizes are too small, higher bin sizes are not architecturally supported and the valid options are automatically populated. Higher values of histogram bins result in greater dynamic range, hence a better resolution at the cost of higher design latency.

Simulation and Results

This example uses an input image of size 240-by-320 pixels, whose path is specified in the inputImage block. The input image pixels are specified by an input bit width of 8 equivalent to `uint8` data type. For 8 tiles in each direction, the computed tile size is 30-by-40 and the number of pixels in each tile is 1200. The number of histogram bins is set to 256.
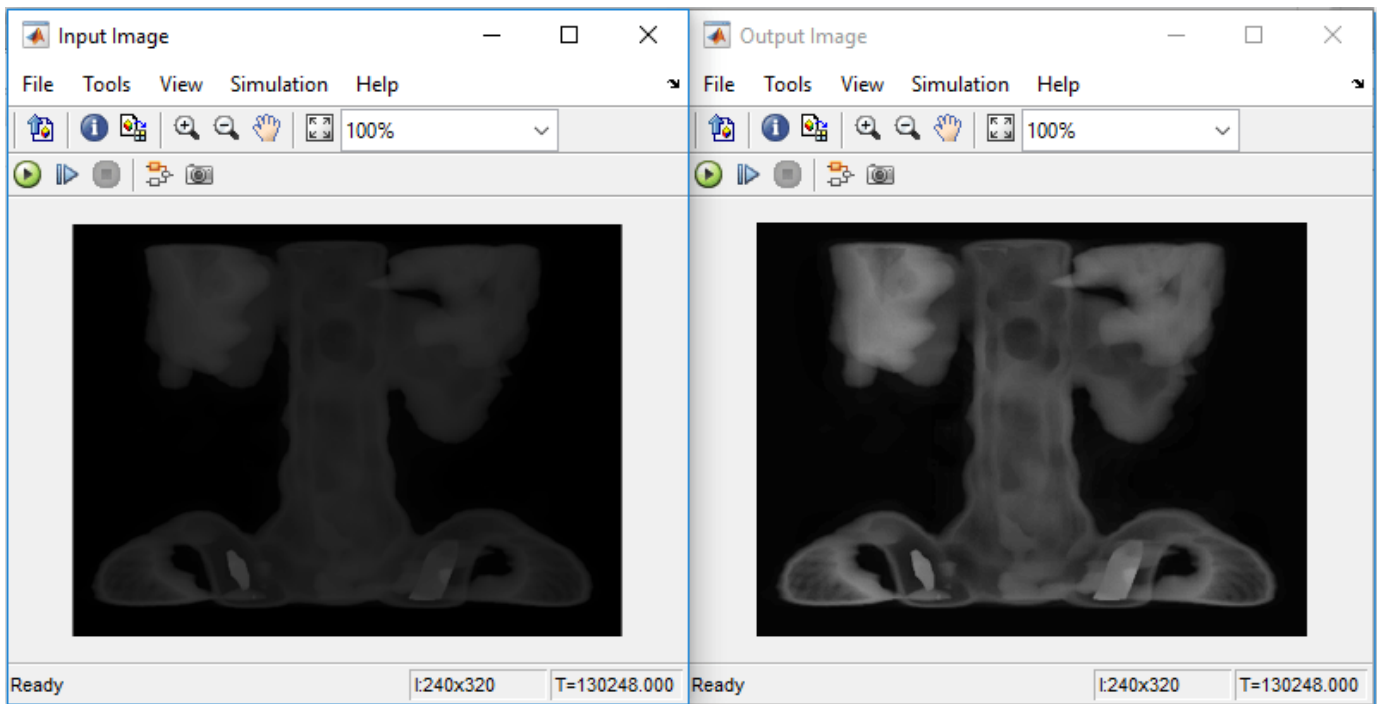
This figure shows the input image and output image from the CLAHE model. The result shows the improved contrast in the output image, without over- saturation. The result of the CLAHE HDL model matches the `adaphisteq` function in MATLAB and has an error of only a few pixels.

HDL code can be generated for the CLAHEHDL subsystem. An HDL Coder™ license is required to generate HDL code. This design was synthesized on the Intel® Arria® 10 GX platform, for 10AX115S2F45I1SG FPGA device. The table shows the resource utilization. The HDL design achieves a clock rate of over 200 MHz.

```
% ================================================
% |Model Name               ||       CLAHEHDL    ||
% ================================================
```

```
% |Input Image Resolution  ||     320 x 240     ||
% |ALM Utilization          ||        48045       ||
% |Total Registers          ||        51688       ||
% |Total RAM Blocks         ||          63        ||
% |Total DSP Blocks         ||           6        ||
% ===============================================
```

References

Karel Zuiderveld, "Contrast Limited Adaptive Histogram Equalization", Graphics Gems IV, p. 474-485, code: p. 479-484.

Change Image Size

This example shows how to downsample a multicomponent image by using the `Image Resizer` block. The example also shows how to implement custom bicubic and Lanczos-2 interpolation algorithms for FPGAs by using basic Simulink blocks.

The most basic of interpolation algorithms, nearest neighbor, assumes the value of its closest neighbor and is computationally lightweight. In theory, you can achieve a more exact reconstruction of an image by using a sinc kernel. However, sinc kernels have infinite spatial extent. To limit the extent, interpolation implementations use simpler kernels to approximate a sinc. The bilinear interpolation algorithm uses the weighted sum of the nearest four pixels to determine the values of the output pixels. Bicubic and Lanczos-2 interpolations are approximations of a sinc kernel. Bicubic interpolation is a more computationally efficient version of the Lanczos-2 method. This example implements and compares these interpolation algorithms.

Behavioral Reference

By default, the `imresize` function uses the bicubic interpolation algorithm. You can choose the nearest neighbor, bilinear, or Lanczos-2 interpolation algorithms by setting the `'Method'` name-value argument to `'nearest'`, `'bilinear'`, or `'lanczos2'`, respectively.

```
v = VideoReader('rhinos.avi');
I = readFrame(v);
Y = imresize(I,[160,256],Method='bilinear');
figure;
imshow(Y)
```


Interpolation Algorithms

Nearest Neighbor

Nearest neighbor interpolation determines the inserted pixel values by assuming the value of the closest of its four neighbors. The interpolator calculates the horizontal, h, and vertical, v, scale factors independently.
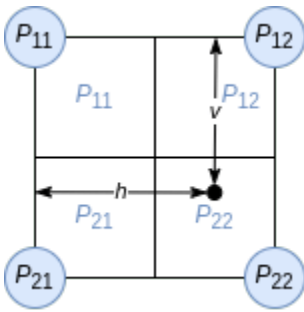
No additional calculations are required once the scale has been determined, which removes the requirement for any multipliers.

Bilinear and Bicubic

Bilinear interpolation, a first-order sinc approximation algorithm, determines the inserted pixel value from the weighted average of the four input pixels nearest to the inserted location.
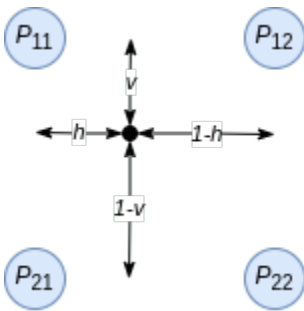
The value for each output pixel is given by $(P_{11}h + P_{12}(1-h))v + (P_{21}h + P_{22}(1-h))(1-v)$.

The bicubic algorithm [1] calculates the average of the 16 input pixels nearest to the inserted location.
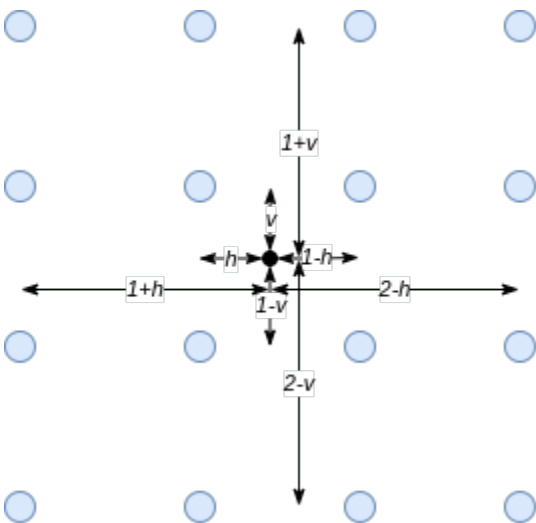
The bicubic coefficients are given by
$$\begin{cases} 1 - 2|d|^2 + |d|^3 & 0 \leq |d| < 1 \\ 4 - 8|d| + 5|d|^2 - |d|^3 & 1 \leq |d| < 2 \\ 0 & 2 \leq |d| \end{cases}.$$

These equations show that the bilinear and bicubic algorithms calculate coefficients for each output pixel.

Lanczos-2

The Lanczos-2 algorithm precalculates the coefficients based on the resize factor. The model calls the `lanczos2_coeffi.m` script to calculate and store these coefficients. The script calculates the Lanczos-2 coefficients using 6 taps and 32 phases.

Implementation of Basic-Blocks Interpolation Algorithms for HDL

This figure shows the principle used to implement the image resize algorithm for hardware. For example, consider resizing an image by a scale factor of 3/4. One possible implementation is to upsample by a factor of 3 and then downsample by a factor of 4. The figure shows the pixel indexes after these operations. Blue dots represent the original pixels, and green crosses represent the interpolated pixels after upsampling.
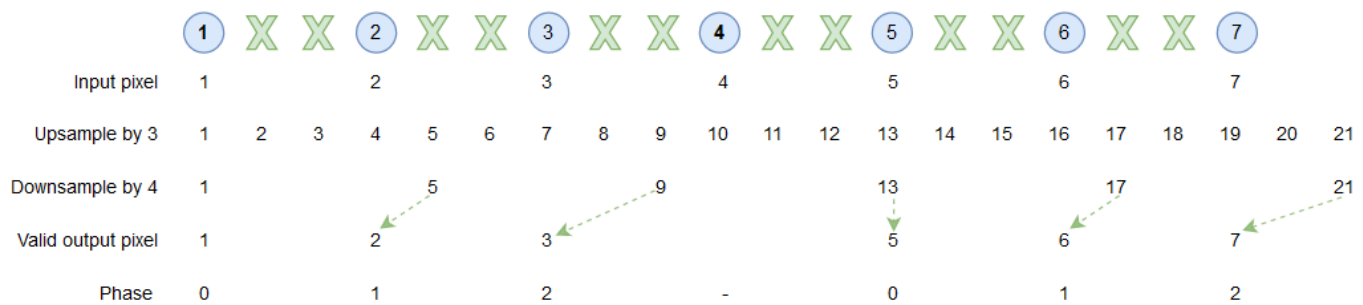
The indexes after downsampling show that not all the interpolated pixels are used in the output image. This example implements a more efficient version of the downsample step by generating interpolated pixels only when they are needed in the output image.

The phase, shown in the bottom line of the figure, is an index that selects which pixels are needed for the output image. When the phase is 0, the algorithm returns the original input pixel value. When the phase is 1, the algorithm calculates coefficients to generate the interpolated pixel in the first position. When the phase is 2, the algorithm calculates coefficients to generate the interpolated pixel in the second position.

Example Model

The model contains two datapaths. The top datapath uses the Image Resizer block to perform resize using bilinear interpolation. The bottom datapath implements resize using bicubic and Lanczos-2 interpolation.
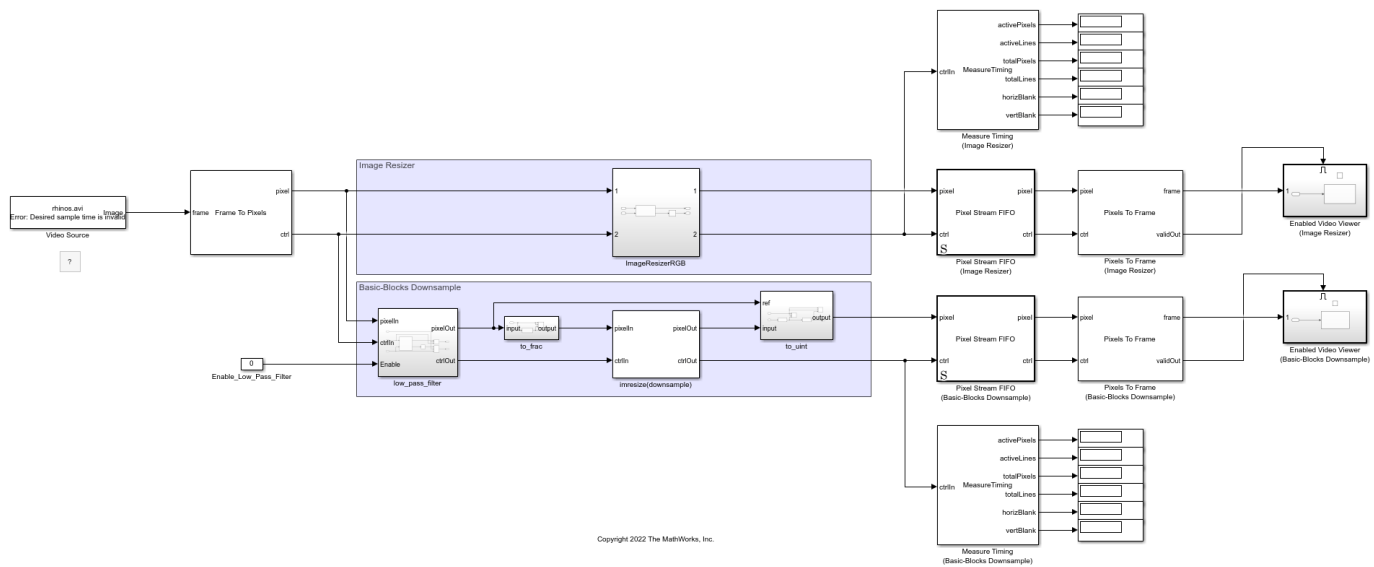
Image Resizer

The `Image Resizer` block provides support for nearest neighbor and bilinear interpolation with an integrated antialiasing prefilter.
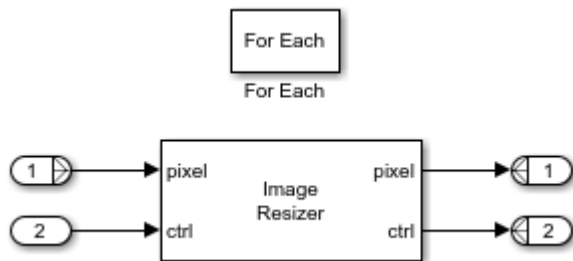
The integrated lowpass Gaussian antialiasing filter provides convenient means to avoid aliasing introduced through the reduction in sampling rate when downsampling.

The `Image Resizer` block provides the ability to insert horizontal blanking between output active video lines. You can use this capability to meet downstream horizontal blanking constraints. However, do not use this capability to perform the role of the `Pixel Stream FIFO`, which buffers input video lines to output contiguous valid active video lines. Inserting too much horizontal blanking can have an impact on the overall frame rate and may necessitate pacing of the input.

The model includes an optional `Pixel Stream FIFO` after the `Image Resizer` block that you can use to consolidate the output pixels, while the `Measure Timing` block displays the size of the output frames.

When processing multicomponent data, use the `Image Resizer` block with a `ForEach` subsystem. The `ForEach` subsystem replicates the `Image Resizer` block across each input component while maintaining a single multicomponent frame input and output.
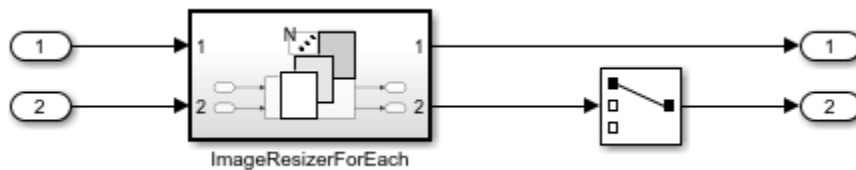
The `ForEach` subsystem is configured to partition and concatenate the pixel input and output on the second dimension, horizontally. The output pixelcontrol buses are all duplicates and so only one needs to be retained at the top level. The `Selector` block selects a single pixelcontrol bus to use in the rest of the data path.

Basic-Blocks Implementation

Similar to the `imresize` function, the `imresize(downsample)` subsystem in this model supports two ways to define the output image size. You can specify a scale factor ranging from 1.000 to 127.999, or you can define the output frame width and height in pixels down to a minimum of 8 by 8. Double-click the `imresize(downsample)` subsystem to set its parameters. The `imresize(downsample)` subsystem requires contiguous video lines.

To avoid aliasing with the basic-blocks implementation, the model includes a lowpass filter before the `imresize(downsample)` subsystem that you can enable by setting its Constant block input to a value of 1. After the `imresize(downsample)` subsystem, there is also an optional `Pixel Stream FIFO` and a `Measure Timing` block to display its output frame sizes.

In the `imresize(downsample)` subsystem, the `input_conversion` and `output_conversion` subsystems convert the color space [2] of the pixel stream based on the parameter on the mask. The `valid_gen_horizontal` and `valid_gen_vertical` subsystems return control signals that are used for generating coefficients and rebuilding the output control bus. If the last line of the image contains no valid pixels after downsampling, the `ctrlBusRebuild` subsystem rebuilds the control bus for the new size.
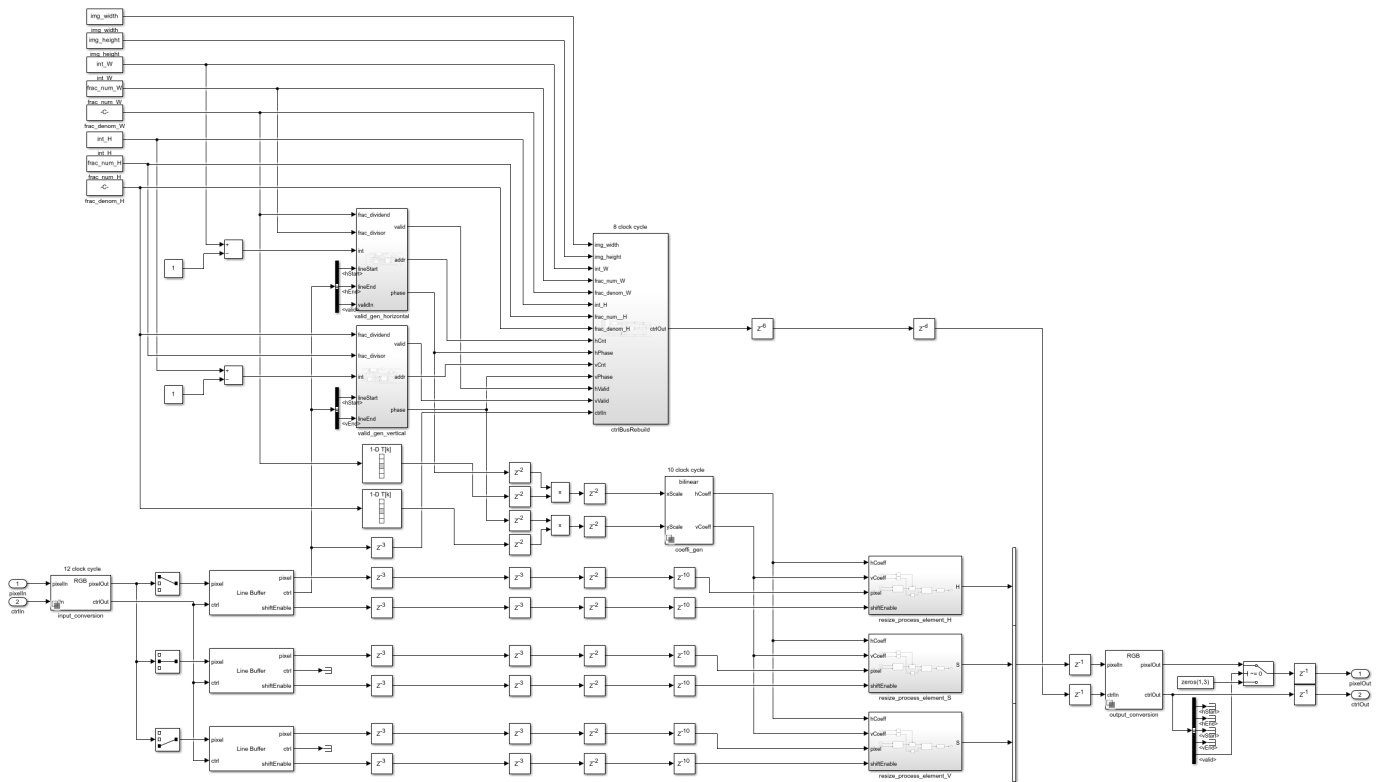
This diagram shows the expected output from the `valid_gen_horizontal` and `valid_gen_vertical` subsystems. The valid signal indicates the validity of the current address and the corresponding phase. To simplify rebuilding the control bus, the first line and row of each output frame are always valid.
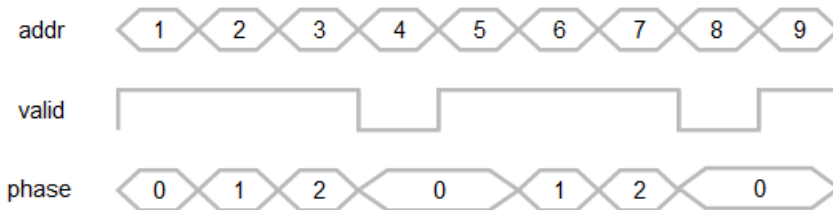
The coefficient generation subsystem, `coeffi_gen`, is a variant subsystem, where bilinear, bicubic, and Lanczos-2 coefficient generators are implemented separately. You can select the algorithm from the mask.
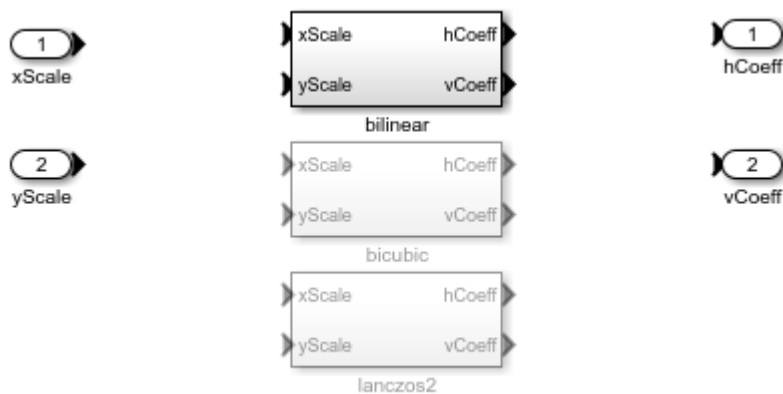
The `resize_process_element` subsystems multiply the coefficients with each pixel component by using a separable filter in vertical order and then in horizontal order. The `trim_0_1` subsystem ensures the result is between 0 and 1.
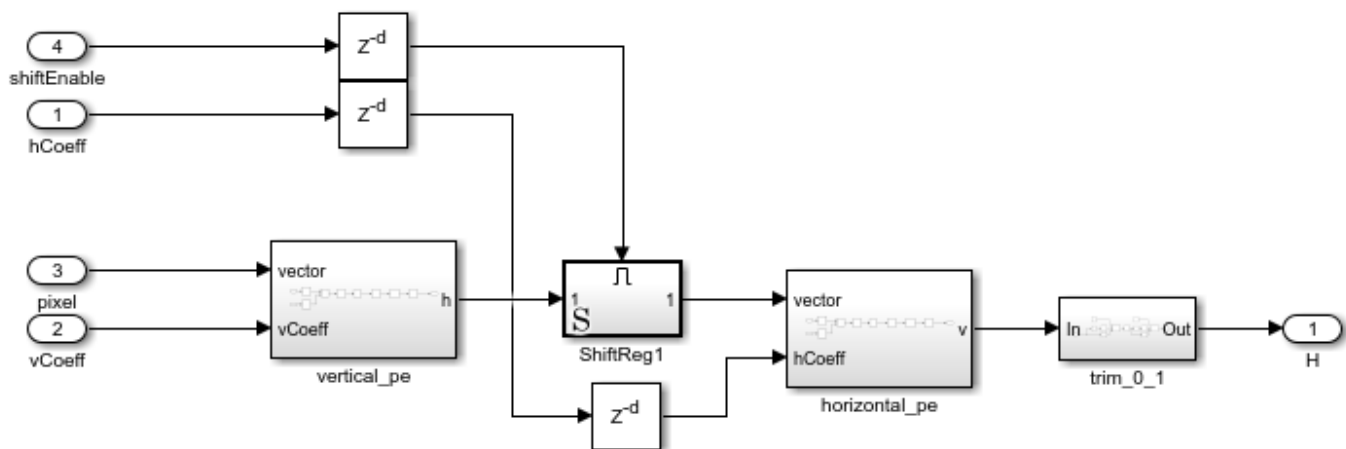
Resource Usage

These tables show the resource usage for the `Image Resizer` block and `imresize(downsample)` subsystem with 240p video input, and do not include the antialiasing filters or the `Pixel Stream FIFOs`. The design was synthesized to achieve a clock frequency of 150 MHz.

This table shows the resources for each of the algorithms when downsampled in the RGB color space.

	LUT	LUTRAM	FF	BRAM	DSP
Image Resizer - Nearest Neighbor	1282	36	1337	3	0
Image Resizer - Bilinear	1317	24	1808	3	12
imresize(downsample) - Bilinear	2414	386	4333	1.5	20
imresize(downsample) - Bicubic	4405	597	8827	4.5	30
imresize(downsample) - Lanczos-2	6490	802	13708	7.5	38

This table shows the resources for each of the three `imresize(downsample)` algorithms when downsampled in the HSV color space.

	LUT	LUTRAM	FF	BRAM	DSP
imresize(downsample) - Bilinear	2987	444	5688	11	26
imresize(downsample) - Bicubic	5117	658	10248	14	36
imresize(downsample) - Lanczos-2	7139	875	15351	17	44

References

[1] Keys, R. "Cubic Convolution Interpolation for Digital Image Processing." *IEEE Transactions on Acoustics, Speech, and Signal Processing* 29, no. 6 (December 1981): 1153-60. https://doi.org/10.1109/TASSP.1981.1163711.

[2] Smith, Alvy Ray. "Color Gamut Transform Pairs." In *Proceedings of the 5th Annual Conference on Computer Graphics and Interactive Techniques - SIGGRAPH '78*, 12-19. Not Known: ACM Press, 1978. https://doi.org/10.1145/800248.807361.

See Also
Image Resizer

Related Examples

Fog Rectification

This example shows how to remove fog from images captured under foggy conditions. The algorithm is suitable for FPGAs.

Fog rectification is an important preprocessing step for applications in autonomous driving and object recognition. Images captured in foggy and hazy conditions have low visibility and poor contrast. These conditions can lead to poor performance of vision algorithms performed on foggy images. Fog rectification improves the quality of the input images to such algorithms.

This example shows a streaming fixed-point implementation of the fog rectification algorithm that is suitable for deployment to hardware.

To improve the foggy input image, the algorithm performs fog removal and then contrast enhancement. The diagram shows the steps of both these operations.

This example takes a foggy RGB image as input. To perform fog removal, the algorithm estimates the dark channel of the image, calculates the airlight map based on the dark channel, and refines the airlight map by using filters. The restoration stage creates a defogged image by subtracting the refined airlight map from the input image.

Then, the Contrast Enhancement stage assesses the range of intensity values in the image and uses contrast stretching to expand the range of values and make features stand out more clearly.
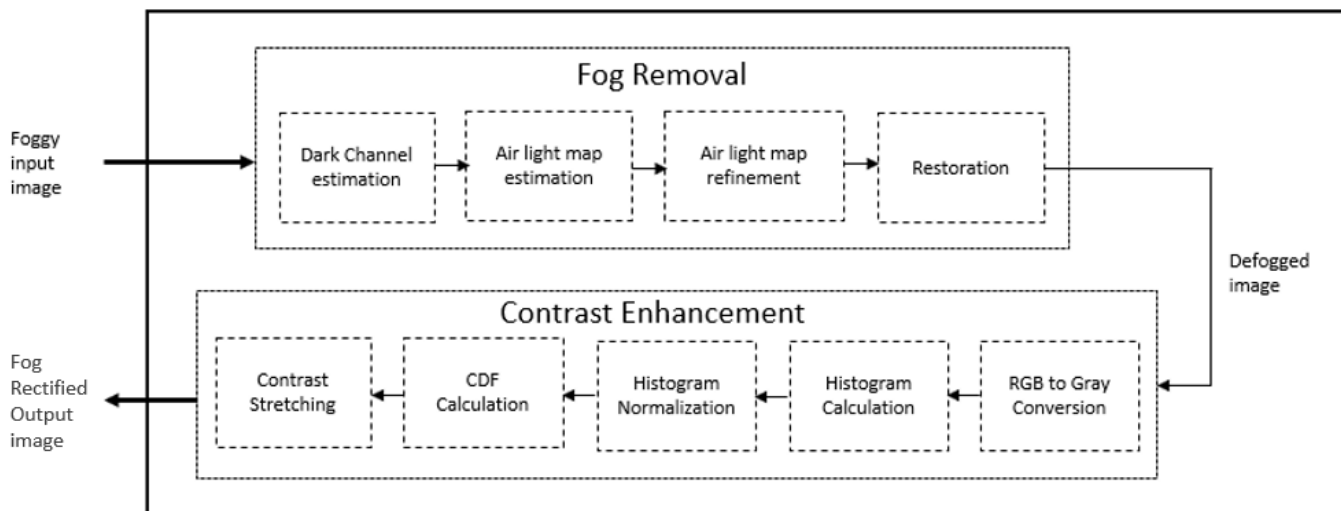
Fog Removal

There are four steps in performing fog removal.

1. **Dark Channel Estimation**: The pixels that represent the non-sky region of an image have low intensities in at least one color component. The channel formed by these low intensities is called the *dark channel*. In a normalized, fog-free image, the intensity of dark channel pixels is very low, nearly zero. In a foggy image, the intensity of dark channel pixels is high, because they are corrupted by fog. So, the fog removal algorithm uses the dark channel pixel intensities to estimate the amount of fog.

The algorithm estimates the dark channel $I_{dark}^c(x,y)$ by finding the pixel-wise minimum across all three components of the input image $I^c(x,y)$ where $c \epsilon [r,g,b]$.

2. **Airlight Map Calculation**: The whiteness effect in an image is known as *airlight*. The algorithm calculates the airlight map from the dark channel estimate by multiplying by a haze factor, z, that represents the amount of haze to be removed. The value of z is between 0 and 1. A higher value means more haze will be removed from the image.

$$I_{air}(x,y) = z \times \min_{c \epsilon [r,g,b]} I_{dark}^c(x,y)$$

3. **Airlight Map Refinement**: The algorithm smoothes the airlight image from the previous stage by using a Bilateral Filter block. This smoothing strengthens the details of the image. The refined image is referred to as $I_{refined}(x,y)$.

4. **Restoration**: To reduce over-smoothing effects, this stage corrects the filtered image using these equations. The constant, m, represents the mid-line of changing the dark regions of the airlight map from dark to bright values. The example uses an empirically derived value of $m = 0.6$.

$$I_{reduced}(x,y) = m \times \ min(I_{air}(x,y), I_{refined}(x,y))$$

The algorithm then subtracts the airlight map from the input foggy image and multiplies by the factor $\dfrac{255}{255 - I_{reduced}(x,y)}$.

$$I_{restore}(x,y) = 255 \times \frac{I^c(x,y) - I_{reduced}(x,y)}{255 - I_{reduced}(x,y)}$$

Contrast Enhancement

There are five steps in contrast enhancement.

1. **RGB to Gray Conversion**: This stage converts the defogged RGB image, $I_{restore}^c(x,y)$, from the fog removal algorithm into a grayscale image, $I_{gray}(x,y)$.

2. **Histogram Calculation**: This stage uses the Histogram block to count the number of pixels falling in each intensity level from 0 to 255.

3. **Histogram Normalization**: The algorithm normalizes the histogram values by dividing them by the input image size.

4. **CDF Calculation**: This stage computes the cumulative distribution function (CDF) of the normalized histogram bin values by adding them to the sum of the previous histogram bin values.

5. **Contrast Stretching**: Contrast stretching is an image enhancement technique that improves the contrast of an image by stretching the range of intensity values to fill the entire dynamic range. When dynamic range is increased, details in the image are more clearly visible.

5a. *i1 and i2 calculation*: This step compares the CDF values with two threshold levels. In this example, the thresholds are 0.05 and 0.95. This calculation determines which pixel intensity values align with the CDF thresholds. These values determine the intensity range for the stretching operation.

5b. *T calculation*: This step calculates the stretched pixel intensity values to meet the desired output intensity values, o_1 and o_2.

o_1 is 10% of maximum output intensity floor(10*255/100) for `uint8` input.

o_2 is 90% of maximum output intensity floor(90*255/100) for `uint8` input.

T is a 256-element vector divided into segments t_1, t_2, and t_3. The segment elements are computed from the relationship between the input intensity range and the desired output intensity range.
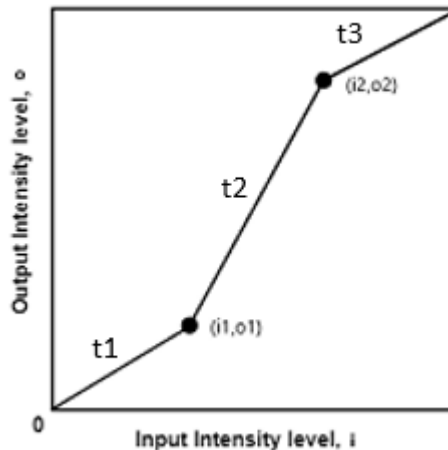
i_1 and i_2 represent two pixel intensities in the input image's range and o_1 and o_2 represent two pixel intensities in the rectified output image's range.

These equations show the how the elements in *T* are calculated.

$$t_1 = \frac{o_1}{i_1}[0:i_1]$$

$$t_2 = (((\frac{o_2 - o_1}{i_2 - i_1})[(i_1 + 1):i_2]) - ((\frac{o_2 - o_1}{i_2 - i_1})i_1)) + o_1$$

$$t_3 = (((\frac{255 - o_2}{255 - i_2})[(i_2 + 1):255]) - ((\frac{255 - o_2}{255 - i_2})i_2)) + o_2$$

$$T = [t_1 \quad t_2 \quad t_3]$$

5c. *Replace intensity values*: This step converts the pixel intensities of the defogged image to the stretched intensity values. Each pixel value in the defogged image is replaced with the corresponding intensity in *T*.

HDL Implementation

The example model implements the algorithm using a steaming pixel format and fixed-point blocks from Vision HDL Toolbox. The serial interface mimics a real time system and is efficient for hardware designs because less memory is required to store pixel data for computation. The serial interface also allows the design to operate independently of image size and format and makes it more resilient to timing errors. Fixed-point data types use fewer resources and give better performance on FPGA. The necessary variables for the example are initialized in the **InitFcn** callback.

Copyright 2019-2022 The MathWorks, Inc.
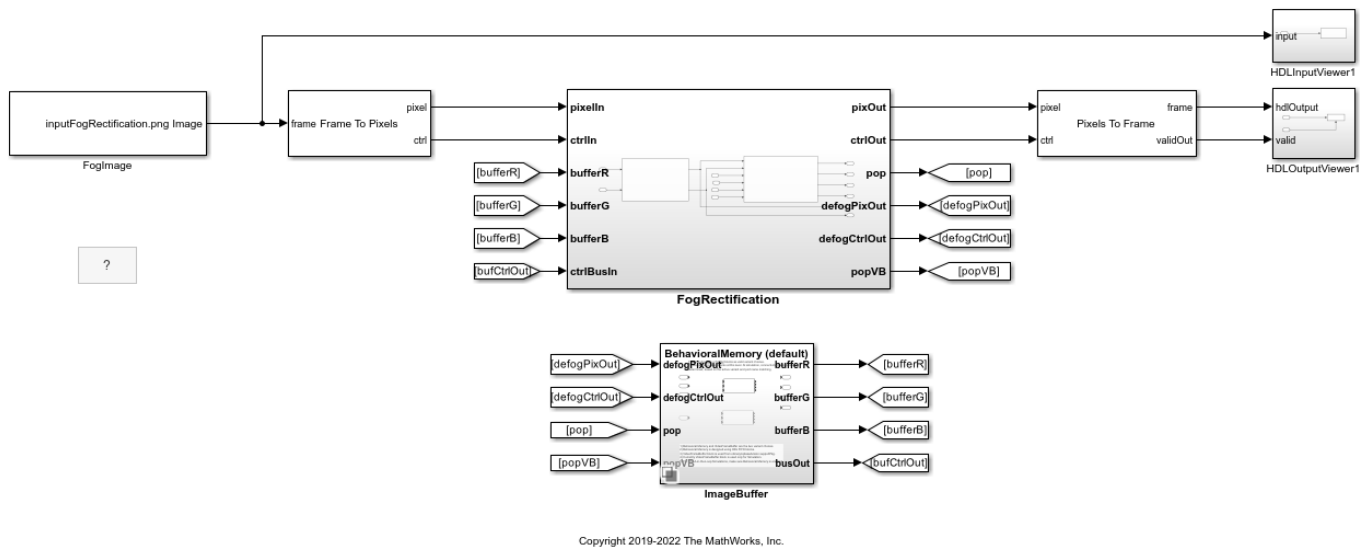
The FogImage block imports the input image to the model. The Frame To Pixels block converts the input frames to a pixel stream of `uint8` values and a `pixelcontrol` bus. The Pixels To Frame block converts the pixel stream back to image frames. The hdlInputViewer subsystem and hdlOutputViewer subsystem show the foggy input image and the defogged enhanced output image, respectively. The ImageBuffer subsystem stores the defogged image so the Contrast Enhancement stages can read it as needed.

The FogRectification subsystem includes the fog removal and contrast enhancement algorithms, implemented with fixed-point datatypes.
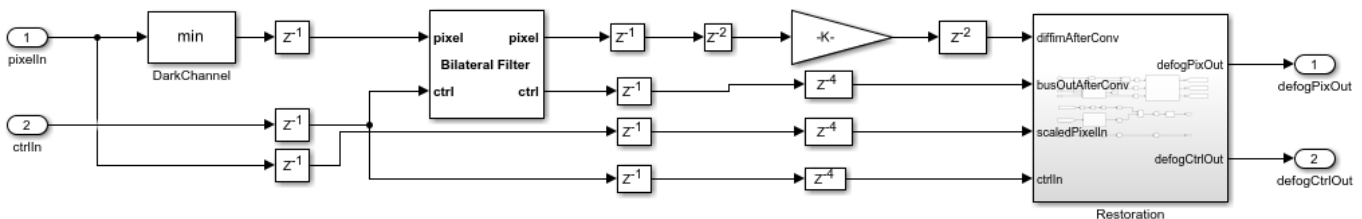
In the FogRemoval subsystem, a Minimum block named DarkChannel calculates the dark channel intensity by finding the minimum across all three components. Then a Bilateral Filter block refines the dark channel results. The filter block has the spatial standard deviation set to 2 and the intensity standard deviation set to 0.5. These parameters are used to derive the filter coefficients. The bit width of the output from filter stage is the same as that of the input.

Next, the airlight image is calculated by multiplying the refined dark channel with a haze factor, 0.9. Multiplying by this factor after the bilateral filter avoids precision loss that would occur from truncating to the maximum 16-bit input size of the bilateral filter.

The Restoration subsystem removes the airlight from the image and then scales the image to prevent over-smoothing. The Pixel Stream Aligner block aligns the input pixel stream with the airlight image before subtraction. The scale factor, m, is found from the midpoint of the difference between the original image and the image with airlight removed. The Restoration subsystem returns a defogged image that has low contrast. So, contrast enhancement must be performed on this image to increase the visibility.

The output from the FogRemoval subsystem is stored in the Image Buffer. The ContrastEnhancement subsystem asserts a `pop` signal to read the frame back from the buffer.
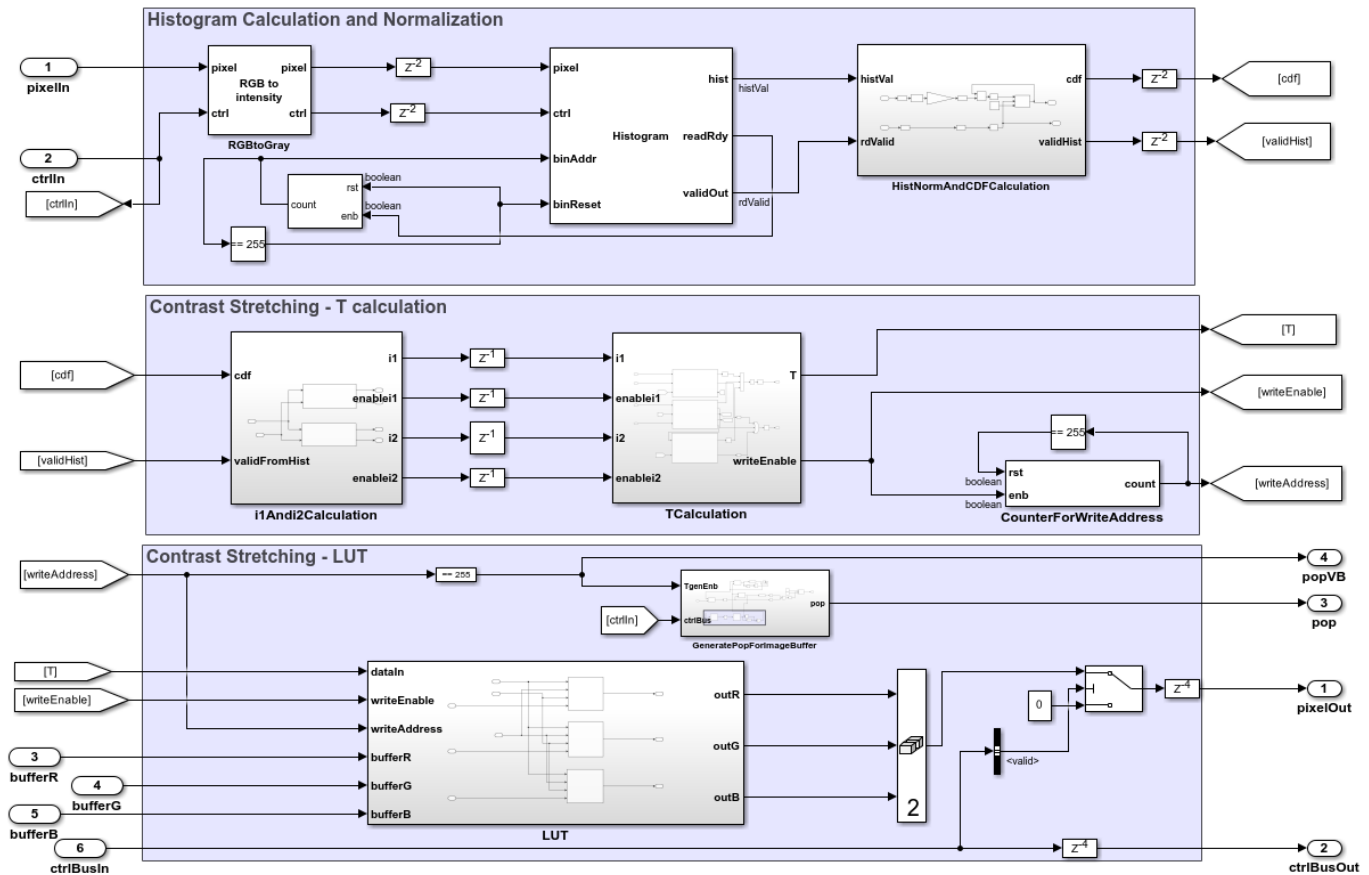
The ContrastEnhancement subsystem uses the Color Space Converter block to convert the RGB defogged image to a grayscale image. Then the Histogram block computes the histogram of pixel intensity values. When the histogram is complete, the block generates a **readRdy** signal. Then the HistNormAndCDFCalculation subsystem normalizes the histogram values and computes the CDF.

The i1Andi2Calculation subsystem computes the i_1 and i_2 values that describe the input intensity range. Then the TCalculation subsystem returns the list of target output intensity values. These 256 values are written into a lookup table. The logic in the Contrast Stretching-LUT area generates a **pop** signal to read the pixel intensities of the defogged image from the Image Buffer, and feeds these values as read addresses to the LUT. The LUT returns the corresponding stretched intensity values defined in *T* to replace the pixel values in the defogged image.

1) Add <u>Subsystem</u> or <u>Model</u> blocks as valid variant choices.
2) You cannot connect blocks at this level. At simulation, connectivity is automatically determined, based on the active variant and port name matching.
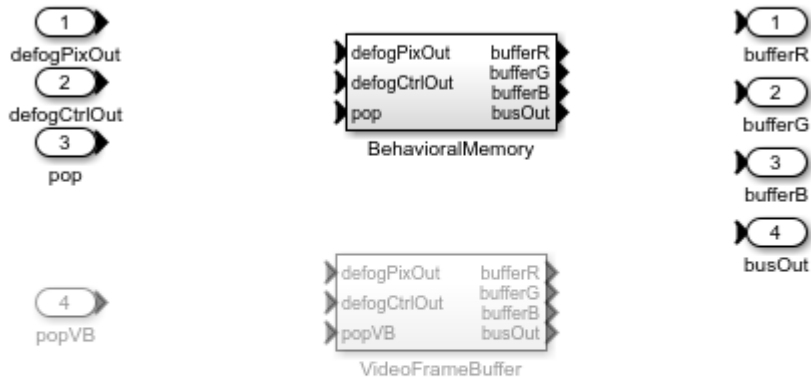
1) Behavioral Memory and VideoFrameBuffer are the two variant choices.
2) Behavioral Memory is designed using HDL FIFO blocks.
3) VideoFrameBuffer block is used from xilinxzynqbasedvision supportPkg.
4) Currently VideoFrameBuffer block is used only for Simulation.
5) For FPGA-in-the-Loop Simulations, make sure Behavioral Memory is enabled.

The Image Buffer subsystem contains two options for modeling the connection to external memory. It is a variant subsystem where you can select between the BehavioralMemory subsystem and the "Model Frame Buffer Interface" (Vision HDL Toolbox Support Package for Xilinx Zynq-Based Hardware) block.
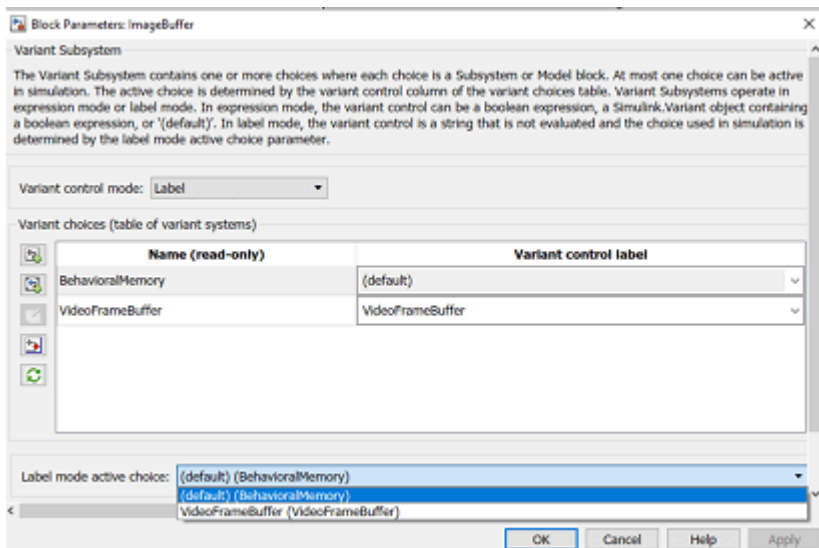
Use the BehavioralMemory subsystem if you do not have the support package mentioned below. This block contains HDL FIFO blocks. The BehavioralMemory returns the stored frame when it receives a pop request signal. The pop request to BehavioralMemory must be high for every row of the frame.

The Video Frame Buffer block requires the Vision HDL Toolbox Support Package for Xilinx Zynq-Based Hardware™. With the proper reference design, the support package can map this block to an AXI-Stream VDMA buffer on the board. This frame buffer returns the stored frame when it receives the popVB request signal. The pop request to this block must be high only one cycle per frame.

The inputs to the Image Buffer subsystem are the pixel stream and control bus generated after fog removal. The pixel stream is fetched during the Contrast Enhancement operation, after the stretched intensities (*T*) are calculated.

Simulation and Results

This example uses an RGB 240-by-320 pixel input image. Both the input pixels and the enhanced output pixels use the `uint8` data type. This design does not have multipixel support.

The figure shows the input and the enhanced output images obtained from the FogRectification subsystem.
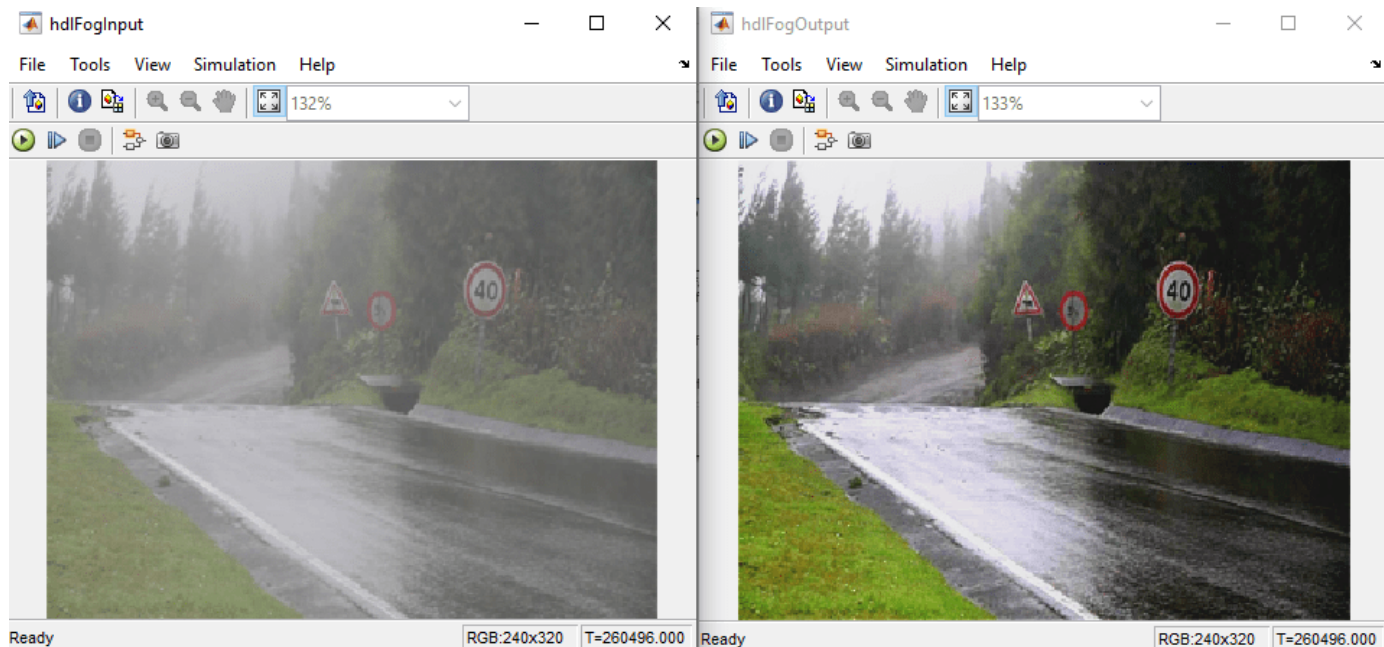
You can generate HDL code for the FogRectification subsystem. An HDL Coder™ license is required to generate HDL code. This design was synthesized for the Intel® Arria® 10 GX (115S2F45I1SG) FPGA. The table shows the resource utilization. The HDL design achieves a clock rate of over 200 MHz.

```
% =================================================================
% |Model Name             ||       FogRectificationHDL     ||
% =================================================================
% |Input Image Resolution ||          320 x 240            ||
% |ALM Utilization        ||            10994              ||
% |Total Registers        ||            20632              ||
% |Total RAM Blocks       ||             67                ||
% |Total DSP Blocks       ||             39                ||
% =================================================================
```

Blob Analysis

This example shows how to implement a single-pass 8-way connected- component labeling algorithm and perform blob analysis.

Blob analysis is a computer vision framework for detection and analysis of connected pixels called blobs. This algorithm can be challenging to implement in a streaming design because it usually involves two or more passes through the image. A first pass performs initial labeling, and additional passes connect any blobs not labeled correctly on the first pass. Streaming designs use a single-pass algorithm to apply and merge labels in hardware and store blob statistics in a RAM. This example has an output stage in software that reads the RAM results and overlays them onto the input video. This example labels blobs, and assigns each blob a unique identifier. Each blob is drawn in a different color in the output image. The example also computes the centroid, bounding box, and area of up to 1024 labeled blobs. The model can support up to 1080p@60 video.

Overview

The example model supports hardware-software co-design. The BlobDetector subsystem is the hardware part of the design, and supports HDL code generation. In a single pass, this subsystem labels each pixel in the incoming pixel stream, merges connected areas, and computes the centroid, area, and bounding box for each blob. The output of the subsystem is a stream of labeled pixels. The subsystem stores the blob statistics in a RAM. When the blob analysis is complete, the subsystem asserts the **data_ready** output port to indicate that the blob statistics are ready to be read.

Logic external to the subsystem reads the statistics one at a time from the BlobDetector RAM by using the **blobIndex** input port as an address. This external logic represents the software part of the design, and does not support HDL code generation. This part of the design reads the centroid, area, and bounding box of each blob, compiles them into vectors for use by the Overlay subsystem, and displays the blob statistics.

The BlobDetector subsystem provides these configuration ports that can be mapped to AXI registers for real-time software control.

- **GradThresh**: Threshold used to create the intensity image.
- **AreaThresh**: Number of pixels that define a blob. The default setting of 1 means that all blobs are processed.
- **CloseOp**: Whether morphological closing is performed prior to labeling and analysis. Closing can be useful after thresholding to fill any introduced holes. By default, this signal is high and enables closing. If you disable closing, the darker coin is detected as two blobs rather than a single connected component.
- **VideoMode**: Pixel stream returned by the subsystem. You can select the input video (0), labeled pixels (1), or intensity video after thresholding (2). You can use these different video views for debugging.

The BlobDetector subsystem returns the output video with associated control signals, and the bounding box, area, and centroid for each requested **blobIndex**. The subsystem also has these output signals to help with debugging.

- **index_o**: Index of the blob currently returning statistics.
- **num_o**: Number of blobs that meet the area threshold.

- **totalNum_o**: Total number of blobs detected in the current frame. By comparing **num_o** and **totalNum_o**, you can fine-tune the input area threshold.
- **data_ready_o**: Indicates when the blob statistics for the current frame are ready to be read from the RAM. In a hardware-software co-design implementation, you can map this signal to an AXI register, and the software can poll the register value to determine when to start reading the statistics.

Copyright 2022 The MathWorks, Inc.
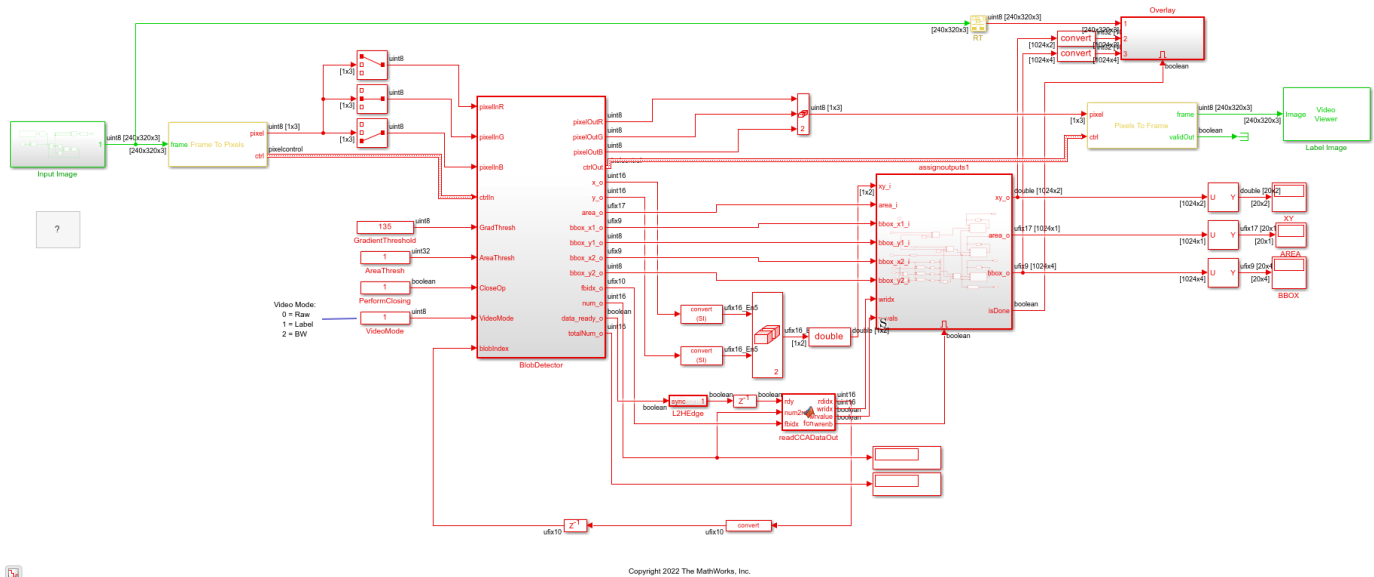
Blob Detector

The BlobDetector subsystem performs connected component labeling and analysis in a single pass over the frame. At the top level, the subsystem contains the CCA_Algorithm subsystem and a cache for the results. The CCA_Algorithm subsystem performs labeling, the calculation of blob statistics, and blob merging.
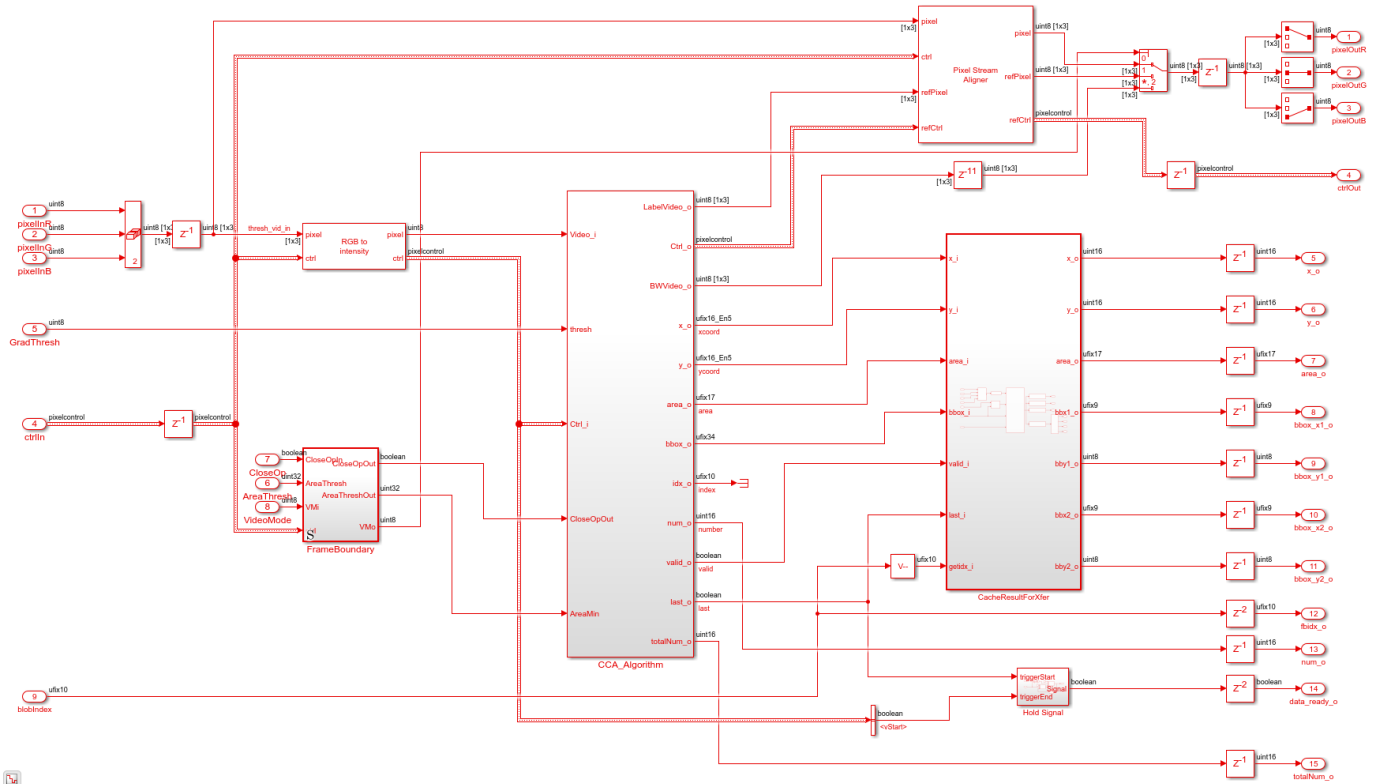
Labeling Algorithm

The labelandmerge MATLAB Function block performs 8-way pixel labeling relative to the current pixel. The possible labels are: previous label, top label, top-left label, and top-right label. The function assigns the current pixel an existing label in order of precedence. If no labels exist, and the pixel is a foreground pixel, then the function assigns a new label to the current pixel by incrementing the label counter. The function forms a labeling window as shown in the diagram by streaming in the current pixel, storing the previous label in a register, and storing the previous line of pixel labels in a RAM. The labels identified by labelandmerge are streamed out of the block as they are identified. For details of the merge operation, see the Merge Logic section.

top-left label	top label	top-right label
previous label	current pixel	
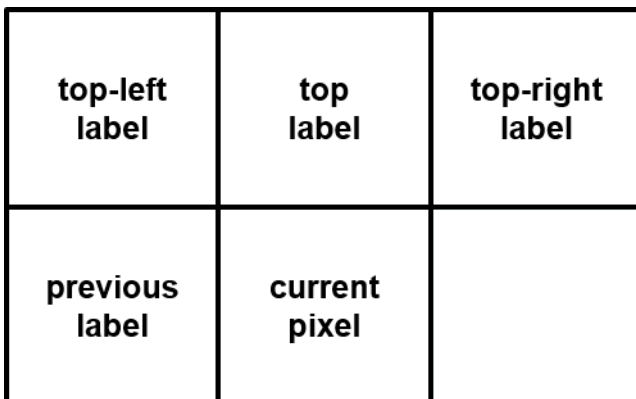
Blob Statistics Calculation

The cca subsystem computes the bounding box, area, and centroid of each blob. This operation uses a set of accumulators and RAMs.

The area_accum subsystem increments the area of the blob represented by each detected label by incrementing a RAM address corresponding to the label.

The x_accum and y_accum subsystems accumulate the **xpos** and **ypos** values from the input ports. The **xpos** and **ypos** values are the coordinates of the pixel in the input frame. Using the area values, and the accumulated coordinates, the centroid is calculated from `xaccum/area` and `yaccum/area`. This calculation uses a single-precision reciprocal for `1/area` and then multiplies that reciprocal by `xaccum` and `yaccum` to find the centroid coordinates. Using a native floating-point reciprocal enables high precision and maintains high dynamic range. When you generate HDL code, the coder implements the reciprocal using fixed-point logic rather than requiring floating-point resources on the FPGA. For more information, see "Getting Started with HDL Coder Native Floating-Point Support" (HDL Coder).

The bbox_store subsystem calculates the bounding box. The subsystem calculates the top-left coordinates, width, and height of the box by comparing the coordinates for each label against the previously cached coordinates.

Merge Logic

During the labeling step, each pixel is examined using only the current line and previous line of label values. This narrow focus means that labels can need correction after further parts of the blob are identified. Label correction can be a challenge for both frame-based and pixel-streaming implementations. The diagrams show two examples of when initial labeling requires correction.

The diagram on the left shows the current pixel connecting two regions through the previous label and top-right label. The diagram on the right shows the current pixel connecting two regions through the previous label and top label. The current pixel is the first location at which the algorithm detects that a merge is required. When the algorithm detects a merge, that pixel is flagged for correction. In both diagrams, the pixels are all part of the same blob and so each pixel must be assigned the same label, 1.
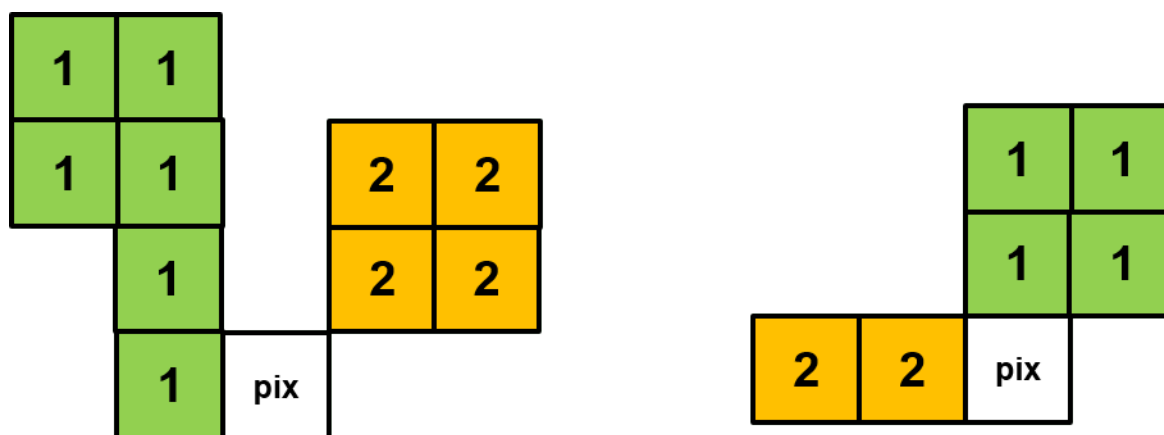
The labelandmerge MATLAB Function block checks for merges and returns a `uint32` value that contains the two merged labels. The MergeQueue subsystem stores any merges that occur on the current line. At the end of each line, the cca subsystem reads the MergeQueue values and corrects

the area, centroid, and bounding box values in the accumulators. The accumulated values for the two merged labels are added together and assigned to a single label. The input to each accumulator subsystem has a 2:1 multiplexer that enables the accumulator to be incremented either when a new label is found, or when a merge occurs.

Output Display

At the end of each frame, the model updates two video displays. The Results On Image video display shows the input image with the bounding boxes (green rectangles) and centroids (red crosses) overlaid. The Label Image video display shows the results of the labeling stage before merging. In the Label Image display, the top of each coin has a different label than the rest of the coin. The merge stage corrects this behavior by merging the two labels into one. The bounding box returned for each blob shows that each coin was detected as a single label.
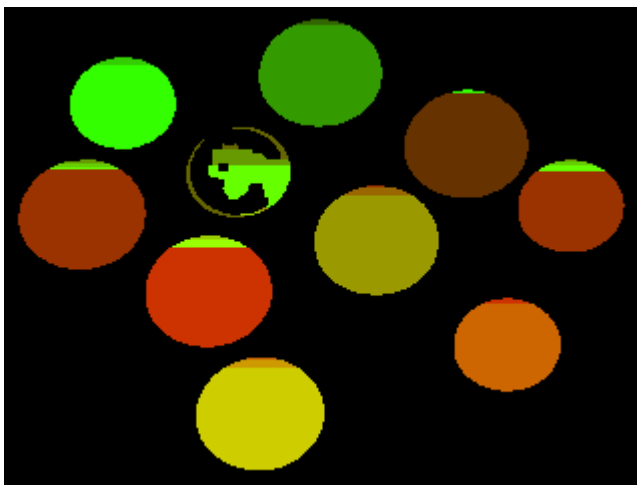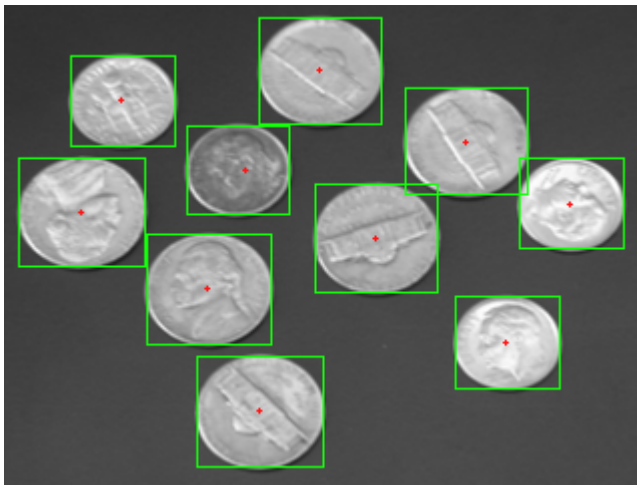
Implementation Results

To check and generate the HDL code referenced in this example, you must have the HDL Coder™ product. To generate the HDL code, use this command.

```
makehdl('BlobAnalysisHDL/BlobDetector')
```

The generated code was synthesized for a target of Xilinx ZC706 SoC. The design met a 200 MHz timing constraint. The design uses very few hardware resources, as shown in the table.

```
T =

  5x2 table

    Resource       Usage
    _____    _____

    DSP48       7 (0.78%)
    Register    4827 (1.1%)
    LUT         3800 (1.74%)
    Slice       1507 (2.67%)
    BRAM        25.5 (4.68%)
```

See Also

More About

- "Hardware-Software Co-Design Workflow for SoC Platforms" (HDL Coder)

Object Tracking using 2-D FFT

This example shows how to implement an object tracking algorithm on FPGA. The model can be configured to support a high frame rate of 1080p@120 fps.

High speed object tracking is essential for a number of computer vision tasks and finds applications ranging across automotive, aerospace and defense sectors. A typical application of such an object tracker could be to precisely guide munitions to hit desired targets. The main principle behind the tracking technique employed is adaptive template matching where the best match of a template within an input image region is detected at each frame.

Download Input File

This example uses the quadrocopter.avi file from the Linkoping Thermal InfraRed (LTIR) dataset [2] as an input. The file is approximately 3 MB in size. Download the file from the MathWorks website and unzip the downloaded file.

```
LTIRZipFile = matlab.internal.examples.downloadSupportFile('visionhdl','LTIR_dataset.zip');
[outputFolder,~,~] = fileparts(LTIRZipFile);
unzip(LTIRZipFile,outputFolder);
quadrocopterVideoFile = fullfile(outputFolder,'LTIR_dataset');
addpath(quadrocopterVideoFile);
```

```
Warning: Cannot overwrite file "C:\Users\nhunting\OneDrive -
MathWorks\Documents\MATLAB\Examples\R2023a\supportfiles\visionhdl\LTIR_dataset\quadrocopter.avi"
The file is already open.
```

Overview

The example model provides two subsystems, a behavioral design using the Computer Vision Toolbox and an HDL design using the Vision HDL Toolbox that is supported for HDL code generation. The ObjectTrackerHDL subsystem is the hardware part of the design, and takes as input a pixel stream. The ROI Selector block dynamically selects an active region of the pixel stream that corresponds to a square search template. This template is 2-D correlated with an initialized adaptive filter. The maximum point of correlation determines the new template location and is used to shift the template in the next frame.

The ObjectTrackerHDL subsystem provides two configuration mask parameters:

- **ObjectCenter**: The x and y coordinate pair that indicates the center of the object or the template.
- **templateSize**: Size of the square template. The allowable sizes range from 16 to 256 taken in the powers of 2.

```
modelname = 'ObjectTrackerHDL';
open_system(modelname);
set_param(modelname,'SampleTimeColors','on');
set_param(modelname,'SimulationCommand','Update');
set_param(modelname,'Open','on');
set(allchild(0),'Visible','off');
```

Copyright 2022-2023 The MathWorks, Inc.
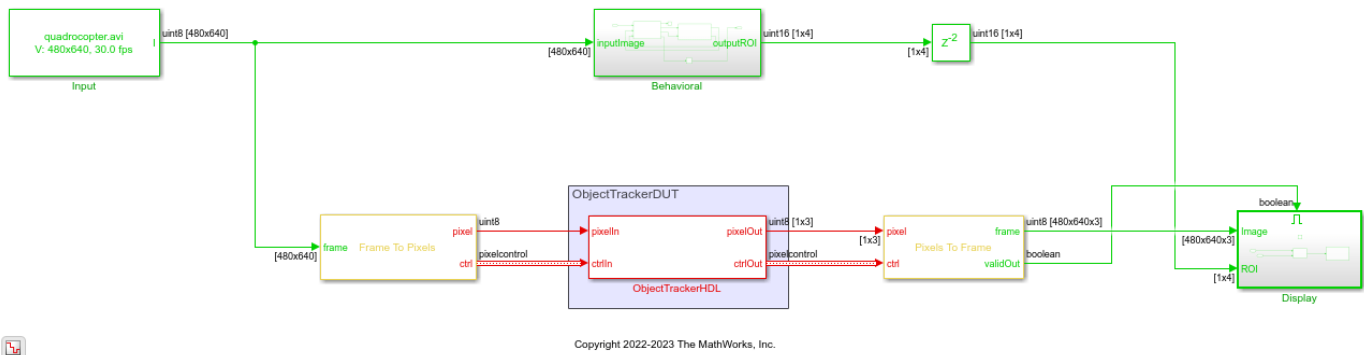
Object Tracker HDL Subsystem

The input to the design is a grayscale or a thermal uint8 image. The input image can be of custom size. Thermal image tracking can involve additional challenges with fast motion and illumination variation. Therefore, a higher frame rate is usually desirable for most InfraRed (IR) applications.

The `ObjectTrackerHDL` design consists of the subsystems: `Preprocess`, `Tracking` and `Overlay` subsystems. The preprocess logic selects the template and does mean subtraction, variance normalization, and windowing to emphasize the target better. Tracking subsystem tracks the template across the frames. The overlay subsystem consists of the `VideoOverlay` block. It accepts a pixel streaming input and takes the position of the template and overlays it onto the frame for viewing. It provides five color options and configurable opacity for better visualization.

```
open_system([modelname '/ObjectTrackerHDL'],'force');
```
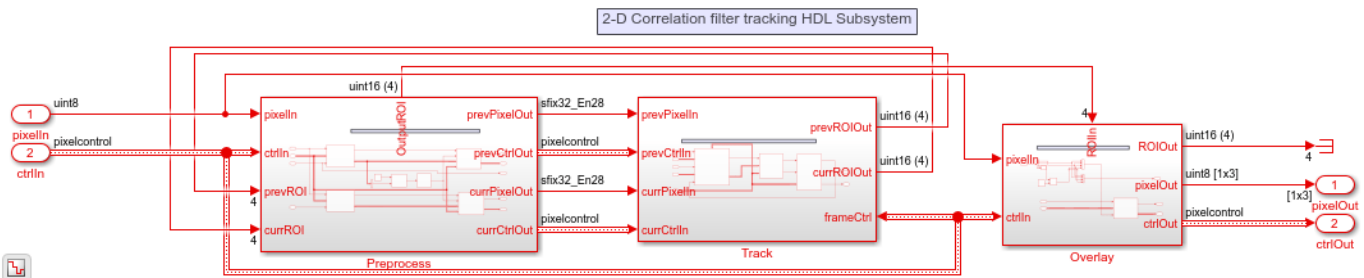

Tracking Algorithm

The tracking algorithm uses a Minimum Output Sum of Squared Error[1] (MOSSE) filter for correlation. This type of filter tries to minimize the sum of squared error between the actual and desired correlation. The initial setup for tracking is a simple training procedure that happens at the initialization of the model. The `InitFcn` callback provides this setup. During this setup, the filter is pre-trained using random affine transformations on the first frame template. The training output is a 2-D Gaussian centered on the training input. The following configurations in the `InitFcn` can be updated additionally to better suit any given application.

- **eta(η)**: The learning rate or the weight given to the previous frame's coefficients.
- **sigma**: The gaussian variance or the sharpness of the target object.
- **trainCount**: The number of training images used.

After the training procedure, the initial coefficients of the filter are available and loaded as constants in the model. This adaptive algorithm updates the filter coefficients after each frame. Let G_i be the

desired correlation output, then the algorithm tries to derive a filter H_i, such that its correlation with the template F_i satisfies the following optimization equation.

$$min \sum_i |F_i \odot H^* - G_i|^2$$

This equation can be solved as follows:

$$H_i = A_i/B_i$$

$$A_i = \eta G_i \odot F_i^* + (1 - \eta)A_{i-1}$$

$$B_i = \eta F_i \odot F_i^* + (1 - \eta)B_{i-1}$$

The learning rate is used to consider the effect of previous frames as the filter adapts to follow the object being tracked. The algorithm is iterative, and the given template is correlated with the filter and the maximum of correlation is used to guide the selection of the new template.

Track Subsystem

After the pixel stream is preprocessed, the `Track` subsystem first performs 2-D correlation between the initial template and the filter. This correlation is performed by first converting the template into frequency domain using 2-D FFT. In the frequency-domain, correlation is efficiently implemented as element-wise multiplication. The `Maxcorrelation` subsystem finds the column and row in the template where the maximum value occurs. It streams in pixels and compares them to find the maximum value and the HV Counter block determines the location of this maximum value. If more than one maximum value exists, it finds the mean of the solutions. If the pixel value streamed in is already equal to the maximum value, the location is updated as the mean location corresponding to both values. This process is repeated until a new maximum value is found or the number of pixels in the frame are exhausted. The `ROIUpdate` subsystem updates the prevROI by using the maximum point in correlation and shifting its center to the new maximum point to yield currROI.

```
open_system([modelname '/ObjectTrackerHDL/Track'],'force');
```
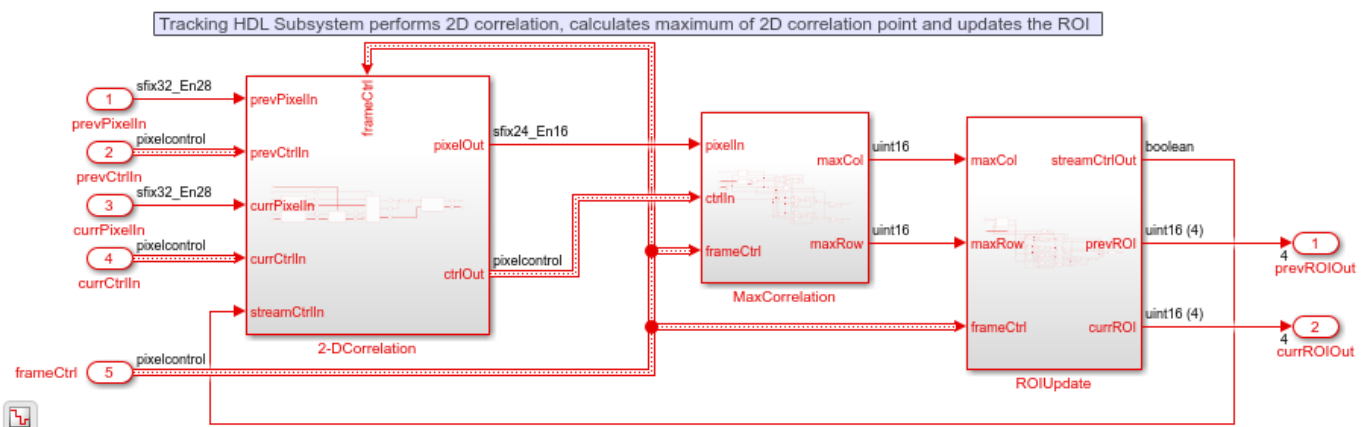

2-D Correlation Subsystem

The 2-D correlation is performed in frequency domain. The `2-DCorrelation` subsystem has two templates in process at each frame, i.e., previous and current ROI templates. Both the templates are represented at the frequency scale using 2-D FFT. The current template is used to update coefficients

of the filter. The `CoefficientsUpdate` subsystem contains RAM blocks to store the coefficients, that are updated to be used in the next frame. The coefficients update block stores the coefficients of the filter in the frequency domain, so they can be element wise multiplied to get the output correlation. The two pixel streams are aligned before multiplication. The alignment is guided by a control determined by comparing the previous and current ROI values. The result is converted back to time domain using an IFFT.

```
open_system([modelname '/ObjectTrackerHDL/Track/2-DCorrelation'],'force');
```
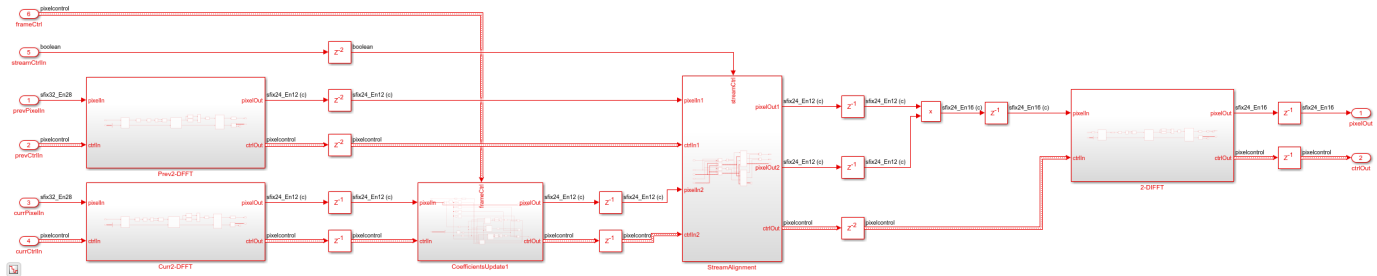

2-D FFT Subsystem

The 2-D FFT is calculated by performing a 1-D FFT across the rows of the template, storing the result and performing a 1-D FFT across its columns. For more details, see FFT (DSP HDL Toolbox)(1-D). The result that is stored in a `CornerTurnMemory` subsystem has ping pong buffering to enable high speed read and write.

```
open_system([modelname '/ObjectTrackerHDL/Track/2-DCorrelation/Prev2-DFFT'],'force');
```
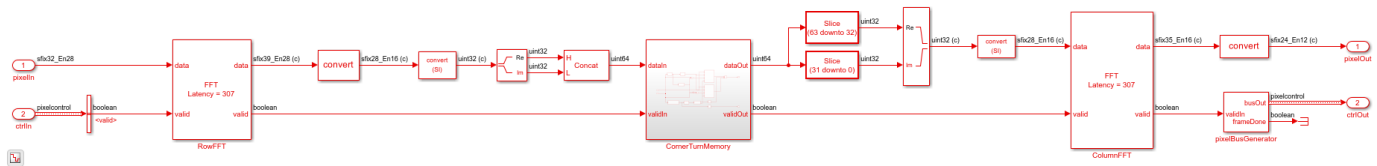

Simulation and Output

At the end of each frame, the model updates the video display for the behavioral and HDL designs. Although the two outputs closely follow each other, a slight deviation in one may compound over a few frames. Both systems can independently track an object through the video. The Linkoping Thermal InfraRed (LTIR) dataset [2] has been used. The quadrocopter sequence is employed in this example and it contains 480p uint8 images and the template size is chosen as 128. The object being tracked is a quadrocopter as shown below.

Implementation Results

To check and generate the HDL code referenced in this example, you must have the HDL Coder™ product. To generate the HDL code, use this command.

```
makehdl('ObjectTrackerHDL/ObjectTrackerHDL')
```

The generated code was synthesized for a target of Xilinx ZCU106 SoC. The design met a 285 MHz timing constraint for a template size of 128. The hardware resources are as shown in the table.

```
T = table(...
    categorical({'DSP48';'Register';'LUT';'BRAM';'URAM'}),...
    categorical({'260 (15.05%)';'65311 (14.17%)';'45414 (19.71%)';'95 (30.44%)';'36 (37.5%)'}),..
    'VariableNames',{'Resource','Usage'})

T =

  5×2 table

    Resource        Usage
    _____    _____

    DSP48       260 (15.05%)
    Register    65311 (14.17%)
    LUT         45414 (19.71%)
    BRAM        95 (30.44%)
    URAM        36 (37.5%)
```

References

[1] D. S. Bolme, J. R. Beveridge, B. A. Draper and Y. M. Lui, "Visual object tracking using adaptive correlation filters," 2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2010, pp. 2544-2550, doi: 10.1109/CVPR.2010.5539960.

[2] A. Berg, J. Ahlberg and M. Felsberg, "A Thermal Object Tracking Benchmark," 2015 12th IEEE International Conference on Advanced Video and Signal Based Surveillance (AVSS), 2015, pp. 1-6, doi: 10.1109/AVSS.2015.7301772.

Ground Plane Segmentation of Lidar Data on FPGA

This example shows how to separate organized 3-D lidar point cloud data into ground and non-ground parts on FPGA. Ground plane removal is an essential preprocessing step in lidar applications.

The example model loads a Velodyne PCAP® file from HDL-32E Lidar sensor. The `PointCloudReader` subsystem uses `readFrame` to read the input lidar data and separate its location and intensity components. The `SegmentGroundFromLidarDataHDL` subsystem performs ground plane segmentation in two major steps, Savitsky - Golay smoothing and breadth-first search ground labeling. The `PointCloudViewer` subsystem creates a point cloud player using `pcplayer` to observe the segmented lidar data. The lidar must be mounted horizontally such that all ground points are observed in the lidar scan closest to the sensor. For more details, see `segmentGroundFromLidarData` (Computer Vision Toolbox).
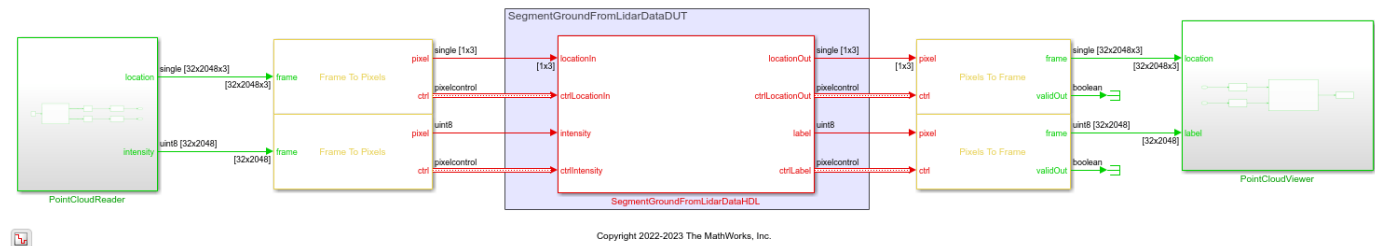
Ground Segmentation HDL Subsystem

The input to the HDL subsystem is an organized point cloud, available as a 32-by-2048-by-3 array. The three channels represent the x-, y-, and z- coordinates of the points. The `CartesianToSphericalProjection` converts these coordinates into a spherical system of range, pitch, and yaw. The `SatizskyGolaySmoothing` subsystem interpolates any non-finite outlier values in the lidar measurements. Finally, the `BFSGroundLabeling` subsystem flood fills the space and labels the connected ground components. To reduce latency and resource use in the flood fill operation, the model reads the values in an upside-down (row-wise) order.
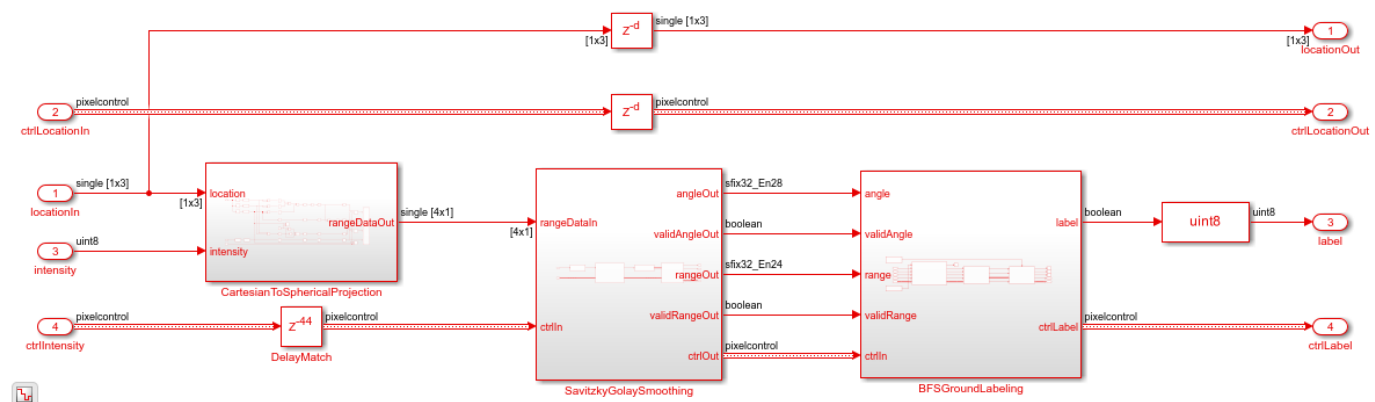
Savitzky Golay Interpolate Subsystem

`SavitzkyGolayInterpolate` subsystem interpolates the non-finite values in the range and pitch plane. For the range plane, the interpolation is done using the neighbors in the column. If the difference of the neighbors and the current pixel is less than the required `rangeThreshold`, the

mean of the differences is the new interpolated value of the current non-finite pixel. For the pitch plane, the interpolation of the current non-finite pixel is done using the left neighbor in the row. Using only the left neighbor optimizes the resources of the design while keeping up with the performance.
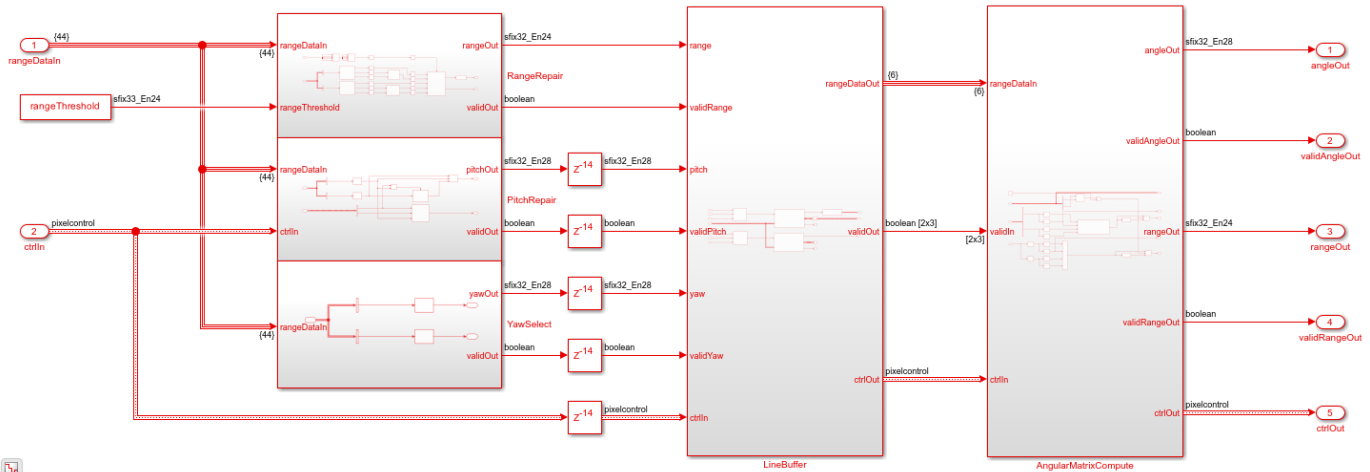
Angle Compute Subsystem

After interpolation, the `AngleCompute` subsystem computes the angle of inclination derived from consecutive range values. It takes each column of the range image and calculates the angle. The `FloodFill` subsystem uses this angle to threshold and label the ground points. Let R_{r-1} and R_r be the range values corresponding to the rows $r-1$ and r in the range image, then the angle can be computed as, $\alpha = atan2(|R_{r-1}\sin(\xi_a) - R_r\sin(\xi_b)|, |R_{r-1}\cos(\xi_a) - R_r\cos(\xi_b)|)$. Here, ξ_a and ξ_b are the vertical angles or pitch values corresponding to the rows $r-1$ and r, respectively.
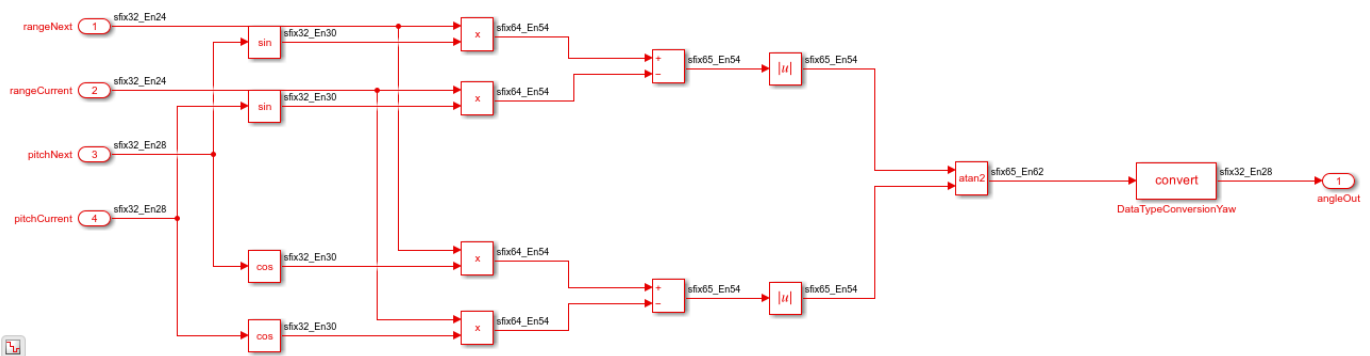
Flood Fill Subsystem

The `BFSGroundLabeling` subsystem uses the flood fill algorithm to segment the ground points. It involves two steps implemented by the subsystems `SeedsCompute` and `FloodFill`. The seeds act as a start point for segmenting the ground points. The `seedThreshold` is the initial elevation angle used to identify the ground point in the scanning line closest to the lidar sensor. `SeedsCompute` subsystem checks if the elevation angle falls below the threshold and marks valid seeds as ground points for every column.

The `FloodFill` subsystem labels points as ground and non-ground points. The flood fill algorithm accuracy increases with more iterations. However, each iteration requires more hardware resources. This example implements a single iteration that filters more than 90% of the points accurately. The `FloodFill` subsystem computes the elevation angle difference between one labeled ground point and its 4-connected neighbors, i.e., top, bottom, left, and right. If any of the four neighbors is a ground pixel and the angle difference is less than a specified `angleThreshold`, then the point is labeled as ground point. The `GroundLabelCompute` subsystem is shown below.
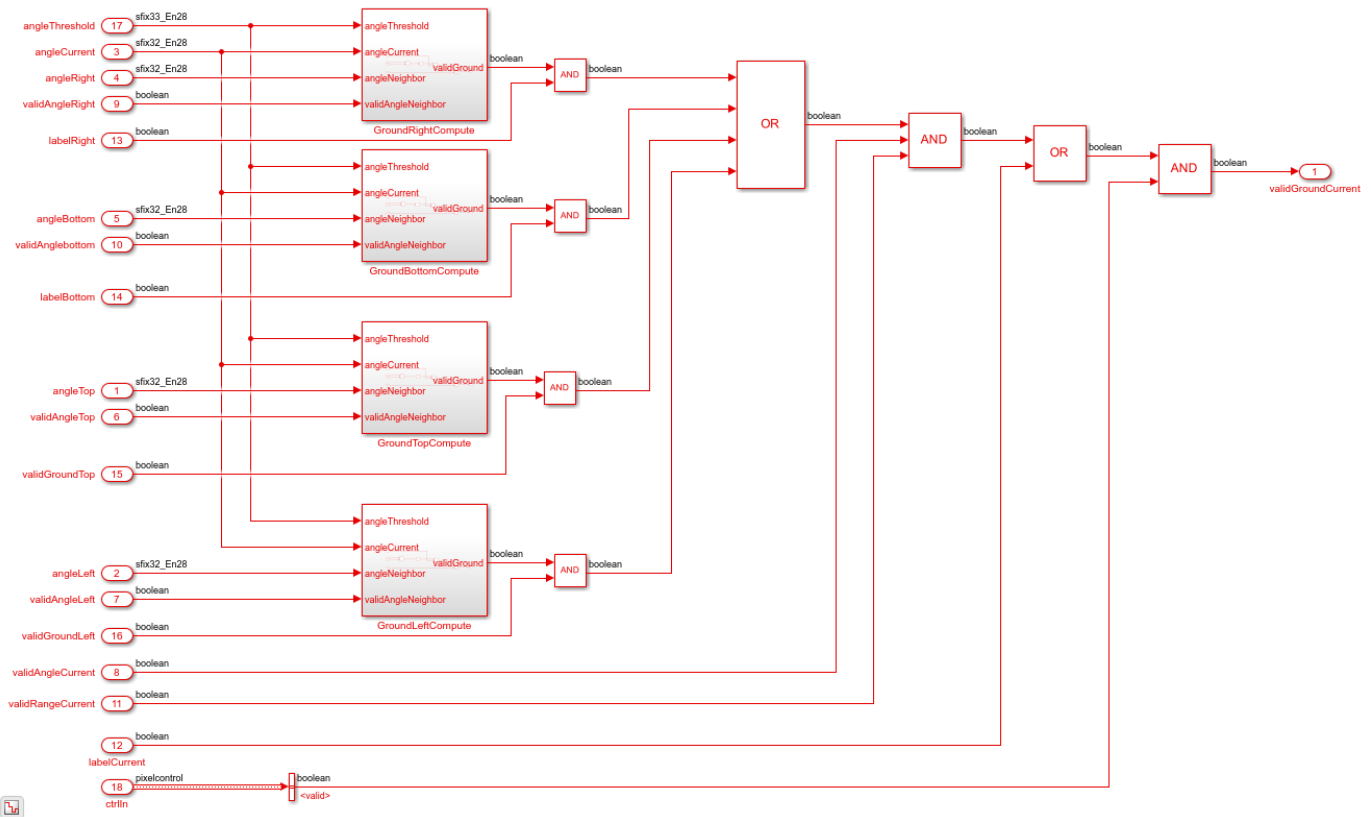
Simulation and Output

The `PointCloudViewer` subsystem uses the ground point labels from the point cloud to color all ground points green and non-ground points white and plot the resulting lidar point cloud. The figure below shows the result of ground plane segmentation for HDL-32E lidar sensor.
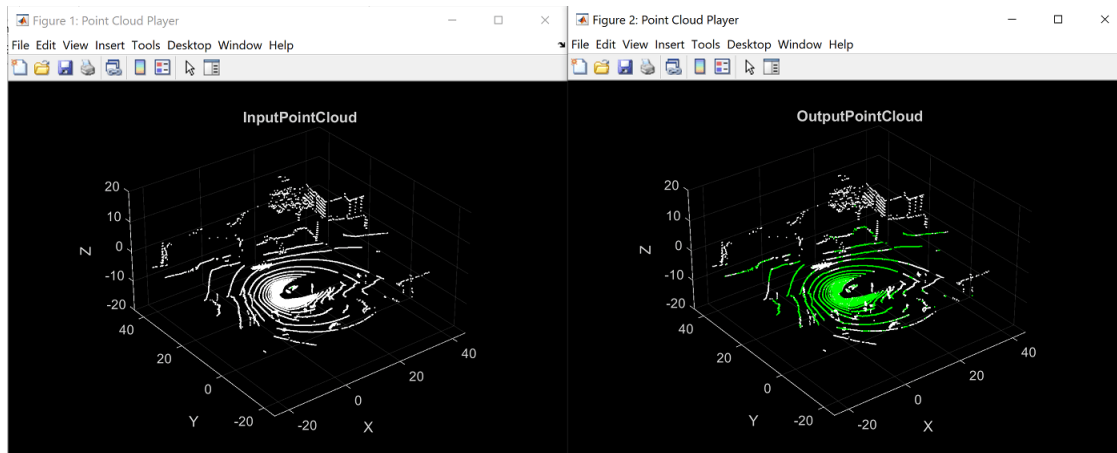
Implementation Results

To check and generate the HDL code referenced in this example, you must have the HDL Coder™ product. To generate the HDL code, use this command.

```
makehdl('SegmentGroundFromLidarDataHDL/SegmentGroundFromLidarDataHDL')
```

The generated code was synthesized for a target of Xilinx ZC706 SoC. The design met a 150 MHz timing constraint. The hardware resources are as shown in the table.

```
T =

  5x2 table

    Resource        Usage
    _____    _____

    DSP48       26 (2.89%)
    Register    39035 (8.93%)
    LUT         33550 (15.35%)
    BRAM        160 (29.36%)
    URAM        0 (0%)
```

References

[1] Bogoslavskyi, I. "Efficient Online Segmentation for Sparse 3D Laser Scans." Journal of Photogrammetry, Remote Sensing and Geoinformation Science. Vol. 85, Number 1, 2017, pp. 41-52.

Pixel-Streaming Design in MATLAB

This example shows how to design pixel-stream video processing algorithms using Vision HDL Toolbox™ objects in the MATLAB® environment and generate HDL code from the design.

This example also tests the design using a small thumbnail image to reduce simulation time. To simulate larger images, such as 1080p video format, use MATLAB Coder™ to accelerate the simulation. See "Accelerate Pixel-Streaming Designs Using MATLAB Coder".

Test Bench

In the test bench `PixelStreamingDesignHDLTestBench.m`, the **videoIn** object reads each frame from a video source which is then converted to grayscale, and then `imresize` is used to reduce this frame from 240p to a thumbnail size for the sake of simulation speed. This thumbnail image is passed to the **frm2pix** object, which converts the full image frame to a stream of pixels and control structures. The function `PixelStreamingDesignHDLDesign.m` is then called to process one pixel (and its associated control structure) at a time. After we process the entire pixel-stream and collect the output stream, the **pix2frm** object converts the output stream to full-frame video. The **viewer** object displays the output and original images side-by-side.

The workflow above is implemented in the following lines of `PixelStreamingDesignHDLTestBench.m`.

```
...
for f = 1:numFrm
    frmFull = rgb2gray(readFrame(videoIn));          % Get a new frame
    frmIn = imresize(frmFull, [actLine actPixPerLine]); % Reduce the frame size

    [pixInVec,ctrlInVec] = frm2pix(frmIn);
    for p = 1:numPixPerFrm
        [pixOutVec(p),ctrlOutVec(p)] = PixelStreamingDesignHDLDesign(pixInVec(p),ctrlInVec(
    end
    frmOut = pix2frm(pixOutVec,ctrlOutVec);

    viewer([frmIn frmOut]);
end
...
```

Both **frm2pix** and **pix2frm** are used to convert between full-frame and pixel-stream domains. The inner for-loop performs pixel-stream processing. The rest of the test bench performs full-frame processing (i.e., **videoIn**, `imresize`, and **viewer**).

Before the test bench terminates, the frame rate is displayed to illustrate the simulation speed.

Pixel-Stream Design

The function defined in `PixelStreamingDesignHDLDesign.m` accepts a pixel stream and five control signals, and returns a modified pixel stream and control signals. For more information on the streaming pixel protocol used by System objects from the Vision HDL Toolbox, see "Streaming Pixel Interface" on page 1-2.

In this example, the function contains the Gamma Corrector System object.

The focus of this example is the workflow, not the algorithm design itself. Therefore, the design code is quite simple. Once you are familiar with the workflow, it is straightforward to implement advanced

video algorithms by taking advantage of the functionality provided by the System objects from Vision HDL Toolbox.

Simulate the Design

Simulate the design with the test bench prior to HDL code generation to make sure there are no runtime errors.

```
PixelStreamingDesignHDLTestBench;
```

```
10 frames have been processed in 35.00 seconds.
Average frame rate is 0.29 frames/second.
```
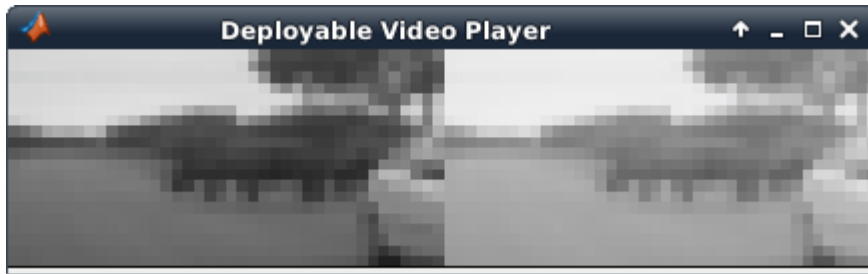

The **viewer** displays the original video on the left, and the output on the right. One can clearly see that the gamma operation results in a brighter image.

Enter the following command to create a new HDL Coder™ project,

```
coder -hdlcoder -new PixelStreamingDesignProject
```

Then, add the file `PixelStreamingDesignHDLDesign.m` to the project as the MATLAB Function and `PixelStreamingDesignHDLTestBench.m` as the MATLAB Test Bench.

Refer to "Get Started with MATLAB to HDL Workflow" (HDL Coder) for a tutorial on creating and populating MATLAB HDL Coder projects.

Launch the Workflow Advisor. In the Workflow Advisor, right-click the 'Code Generation' step. Choose the option 'Run to selected task' to run all the steps from the beginning through HDL code generation.

Examine the generated HDL code by clicking the links in the log window.

Enhanced Edge Detection from Noisy Color Video

This example shows how to develop a complex pixel-stream video processing algorithm, accelerate its simulation using MATLAB® Coder™, and generate HDL code from the design. The algorithm enhances the edge detection from noisy color video.

You must have a MATLAB Coder license to run this example.

This example builds on the "Pixel-Streaming Design in MATLAB" on page 2-206 and the "Accelerate Pixel-Streaming Designs Using MATLAB Coder" examples.

Test Bench

In the `EnhancedEdgeDetectionHDLTestBench.m` file, the **videoIn** object reads each frame from a color video source, and the `imnoise` function adds salt and pepper noise. This noisy color image is passed to the **frm2pix** object, which converts the full image frame to a stream of pixels and control structures. The function `EnhancedEdgeDetectionHDLDesign.m` is then called to process one pixel (and its associated control structure) at a time. After we process the entire pixel-stream and collect the output stream, the **pix2frm** object converts the output stream to full-frame video. A full-frame reference design `EnhancedEdgeDetectionHDLReference.m` is also called to process the noisy color image. Its output is compared with that of the pixel-stream design. The function `EnhancedEdgeDetectionHDLViewer.m` is called to display video outputs.

The workflow above is implemented in the following lines of `EnhancedEdgeDetectionHDLTestBench.m`.

```
...
frmIn = zeros(actLine,actPixPerLine,3,'uint8');
for f = 1:numFrm
    frmFull = readFrame(videoIn);              % Get a new frame
    frmIn = imnoise(frmFull,'salt & pepper'); % Add noise

    % Call the pixel-stream design
    [pixInVec,ctrlInVec] = frm2pix(frmIn);
    for p = 1:numPixPerFrm
        [pixOutVec(p),ctrlOutVec(p)] = EnhancedEdgeDetectionHDLDesign(pixInVec(p,:),ctrlInVec(p)
    end
    frmOut = pix2frm(pixOutVec,ctrlOutVec);

    % Call the full-frame reference design
    [frmGray,frmDenoise,frmEdge,frmRef] = EnhancedEdgeDetectionHDLReference(frmIn);

    % Compare the results
    if nnz(imabsdiff(frmRef,frmOut))>20
        fprintf('frame %d: reference and design output differ in more than 20 pixels.\n',f);
        return;
    end

    % Display the results
    EnhancedEdgeDetectionHDLViewer(actPixPerLine,actLine,[frmGray frmDenoise uint8(255*[frmEdge
end
...
```

Since frmGray and frmDenoise are uint8 data type while frmEdge and frmOut are logical, **uint8(255x[frmEdge frmOut])** maps logical false and true to uint8(0) and uint8(255), respectively, so that matrices can be concatenated.

Both **frm2pix** and **pix2frm** are used to convert between full-frame and pixel-stream domains. The inner for-loop performs pixel-stream processing. The rest of the test bench performs full-frame processing.

Before the test bench terminates, frame rate is displayed to illustrate the simulation speed.

For the functions that do not support C code generation, such as `tic`, `toc`, `imnoise`, and `fprintf` in this example, use **coder.extrinsic** to declare them as extrinsic functions. Extrinsic functions are excluded from MEX generation. The simulation executes them in the regular interpreted mode. Since `imnoise` is not included in the C code generation process, the compiler cannot infer the data type and size of frmIn. To fill in this missing piece, we add the statement **frmIn = zeros(actLine,actPixPerLine,3,'uint8')** before the outer for-loop.

Pixel-Stream Design

The function defined in `EnhancedEdgeDetectionHDLDesign.m` accepts a pixel stream and a structure consisting of five control signals, and returns a modified pixel stream and control structure. For more information on the streaming pixel protocol used by System objects from the Vision HDL Toolbox, see the "Streaming Pixel Interface" on page 1-2.

In this example, the **rgb2gray** object converts a color image to grayscale, **medfil** removes the salt and pepper noise. **sobel** highlights the edge. Finally, the **mclose** object performs morphological closing to enhance the edge output. The code is shown below.

```
[pixGray,ctrlGray] = rgb2gray(pixIn,ctrlIn);        % Convert RGB to grayscale
[pixDenoise,ctrlDenoise] = medfil(pixGray,ctrlGray); % Remove noise
[pixEdge,ctrlEdge] = sobel(pixDenoise,ctrlDenoise);  % Detect edges
[pixClose,ctrlClose] = mclose(pixEdge,ctrlEdge);     % Apply closing
```

Full-Frame Reference Design

When designing a complex pixel-stream video processing algorithm, it is a good practice to develop a parallel reference design using functions from the Image Processing Toolbox™. These functions process full image frames. Such a reference design helps verify the implementation of the pixel-stream design by comparing the output image from the full-frame reference design to the output of the pixel-stream design.

The function `EnhancedEdgeDetectionHDLReference.m` contains a similar set of four functions as in the `EnhancedEdgeDetectionHDLDesign.m`. The key difference is that the functions from Image Processing Toolbox process full-frame data.

Due to the implementation difference between `edge` function and `visionhdl.EdgeDetector` System object, reference and design output are considered matching if frmOut and frmRef differ in no greater than 20 pixels.

Create MEX File and Simulate the Design

Generate and execute the MEX file.

```
codegen('EnhancedEdgeDetectionHDLTestBench');
```

Code generation successful.

```
EnhancedEdgeDetectionHDLTestBench_mex;
```

frame 1: reference and design output differ in more than 20 pixels.
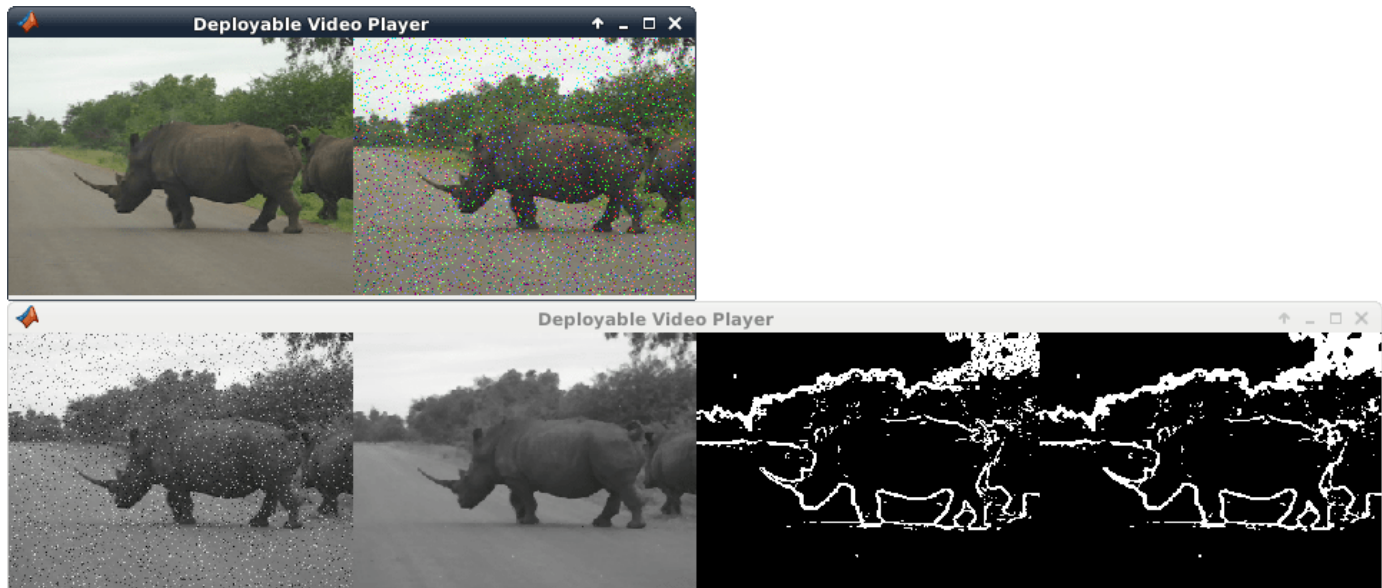
The upper video player displays the original color video on the left, and its noisy version after adding salt and pepper noise on the right. The lower video player, from left to right, represents: the grayscale image after color space conversion, the de-noised version after median filter, the edge output after edge detection, and the enhanced edge output after morphological closing operation.

Note that in the lower video chain, only the enhanced edge output (right-most video) is generated from pixel-stream design. The other three are the intermediate videos from the full-frame reference design. To display all of the four videos from the pixel-stream design, you would have written the design file to output four sets of pixels and control signals, and instantiated three more **visionhdl.PixelsToFrame** objects to convert the three intermediate pixel streams back to frames. For the sake of simulation speed and the clarity of the code, this example does not implement the intermediate pixel-stream displays.

HDL Code Generation

To create a new project, enter the following command in the temporary folder

```
coder -hdlcoder -new EnhancedEdgeDetectionProject
```

Then, add the file 'EnhancedEdgeDetectionHDLDesign.m' to the project as the MATLAB Function and 'EnhancedEdgeDetectionHDLTestBench.m' as the MATLAB Test Bench.

Refer to "Get Started with MATLAB to HDL Workflow" (HDL Coder) for a tutorial on creating and populating MATLAB HDL Coder projects.

Launch the Workflow Advisor. In the Workflow Advisor, right-click the 'Code Generation' step. Choose the option 'Run to selected task' to run all the steps from the beginning through HDL code generation.

Examine the generated HDL code by clicking the links in the log window.

Code Generation and Deployment

Accelerate a MATLAB Design with MATLAB Coder

Vision HDL Toolbox designs in MATLAB must call one or more System objects for every pixel. This serial processing is efficient in hardware, but is slow in simulation. One way to accelerate simulations of these objects is to simulate using generated C code rather than the MATLAB interpreted language.

Code generation accelerates simulation by using constants for the sizes and data types of variables inside the function. This process removes the overhead of the interpreted language checking for size and data type in every line of code. You can compile a video processing algorithm and test bench into MEX functions, and use the resulting MEX file to speed up the simulation.

To generate C code, you must have a MATLAB Coder™ license.

See "Accelerate Pixel-Streaming Designs Using MATLAB Coder".

HDL Code Generation from Vision HDL Toolbox

In this section...
"What Is HDL Code Generation?" on page 3-3
"HDL Code Generation Support in Vision HDL Toolbox" on page 3-3
"Streaming Pixel Interface in HDL" on page 3-3

What Is HDL Code Generation?

You can use MATLAB and Simulink for rapid prototyping of hardware designs. Vision HDL Toolbox blocks and System objects, when used with HDL Coder™, provide support for HDL code generation. HDL Coder tools generate target-independent synthesizable Verilog® and VHDL® code for FPGA programming or ASIC prototyping and design.

HDL Code Generation Support in Vision HDL Toolbox

Most blocks and objects in Vision HDL Toolbox support HDL code generation.

The following blocks and objects are for simulation only and are not supported for HDL code generation :

- Frame To Pixels (`visionhdl.FrameToPixels`)
- Pixels To Frame (`visionhdl.PixelsToFrame`)
- FIL Frame To Pixels (`visionhdl.FILFrameToPixels`)
- FIL Pixels To Frame (`visionhdl.FILPixelsToFrame`)
- Measure Timing (`visionhdl.MeasureTiming`)

Streaming Pixel Interface in HDL

The streaming pixel bus and structure data type used by Vision HDL Toolbox blocks and System objects is flattened into separate signals in HDL.

In VHDL, the interface is declared as:

```
PORT( clk              :   IN    std_logic;
      reset            :   IN    std_logic;
      enb              :   IN    std_logic;
      in0              :   IN    std_logic_vector(7 DOWNTO 0); -- uint8
      in1_hStart       :   IN    std_logic;
      in1_hEnd         :   IN    std_logic;
      in1_vStart       :   IN    std_logic;
      in1_vEnd         :   IN    std_logic;
      in1_valid        :   IN    std_logic;
      out0             :   OUT   std_logic_vector(7 DOWNTO 0); -- uint8
      out1_hStart      :   OUT   std_logic;
      out1_hEnd        :   OUT   std_logic;
      out1_vStart      :   OUT   std_logic;
      out1_vEnd        :   OUT   std_logic;
```

```
        out1_valid      :   OUT   std_logic
        );
```

In Verilog, the interface is declared as:

```
input   clk;
input   reset;
input   enb;
input   [7:0] in0;  // uint8
input   in1_hStart;
input   in1_hEnd;
input   in1_vStart;
input   in1_vEnd;
input   in1_valid;
output  [7:0] out0;  // uint8
output  out1_hStart;
output  out1_hEnd;
output  out1_vStart;
output  out1_vEnd;
output  out1_valid;
```

Blocks and System Objects Supporting HDL Code Generation

Most blocks and objects in Vision HDL Toolbox are supported for HDL code generation. For exceptions, see "HDL Code Generation Support in Vision HDL Toolbox" on page 3-3. This page helps you find blocks and objects supported for HDL code generation in other MathWorks® products.

Blocks

To create a library of HDL-supported blocks from all your installed products, enter `hdllib` at the MATLAB command line. This command requires an HDL Coder license.

You can also view blocks that are supported for HDL code generation in documentation by filtering the block reference list. Click **Blocks** in the blue bar at the top of the Help window, then select the **HDL code generation** check box at the bottom of the left column. The blocks are listed in their respective products. You can use the table of contents in the left column to navigate between products and categories.

Refer to the "Extended Capabilities > HDL Code Generation" section of each block page for block implementations, properties, and restrictions for HDL code generation.

System Objects

You can view System objects that are supported for HDL code generation in documentation by filtering the functions reference list. Click **Functions** in the blue bar at the top of the Help window, then select the **HDL code generation** check box at the bottom of the left column. The System objects are listed in their respective products. You can use the table of contents in the left column to navigate between products and categories.

Refer to the "Extended Capabilities > HDL Code Generation" section of each block page for restrictions for HDL code generation.

Generate HDL Code from Simulink

Introduction

This page shows you how to generate HDL code from the design described in "Design Video Processing Algorithms for HDL in Simulink". You can generate HDL code from the HDL Algorithm subsystem in the model.

To generate HDL code, you must have an HDL Coder license.

Prepare Model

Run the `visionhdlsetup` function to configure the model for HDL code generation. If you started your design using the Vision HDL Toolbox Simulink model template, your model is already configured for HDL code generation.

Generate HDL Code

Right-click the HDL Algorithm block, and select **HDL Code > Generate HDL from subsystem** to generate HDL using the default settings. The output log of this operation is shown in the MATLAB Command Window, along with the location of the generated files.

To change code generation options, use the **HDL Code Generation** section of Simulink Configuration Parameters. For guidance through the HDL code generation process, or to select a target device or synthesis tool, right-click on the HDL Algorithm block, and select **HDL Code > HDL Workflow Advisor**.

Alternatively, from the MATLAB Command Window, you can call:

```
makehdl([modelname '/HDL Algorithm'])
```

Generate HDL Test Bench

You can select options to generate a test bench in Simulink Configuration Parameters or in **HDL Workflow Advisor**.

Alternatively, to generate an HDL test bench from the command line, call:

```
makehdltb([modelname '/HDL Algorithm'])
```

See Also

Functions
makehdl | makehdltb

Related Examples

- "HDL Code Generation and FPGA Synthesis from Simulink Model" (HDL Coder)
- "Choose a Test Bench for Generated HDL Code" (HDL Coder)

Generate HDL Code from MATLAB

This page shows you how to generate HDL code from the design in the "Design Hardware-Targeted Image Filters in MATLAB" example.

To generate HDL code, you must have an HDL Coder license.

Create an HDL Coder Project

Start from a working folder that contains a function and a test bench MATLAB script file. You can use this `openExample` command to get a working folder with the files from the "Design Hardware-Targeted Image Filters in MATLAB" example.

```
openExample('visionhdl/VisionHDLMATLABTutorialExample')
```

Open the HDL Coder app and create a new project.

```
coder -hdlcoder -new vht_matlabhdl_ex
```

In the **HDL Code Generation** pane, add the function file `HDLTargetedDesign.m` and the test bench file `VisionHDLMATLABTutorialExample.m` to the project.

Click next to the signal names under **MATLAB Function** to define the data types for the input and output signals of the function. The control signals are a `struct` of five `logical` scalars, enter their names as shown in the figure. The pixel data type is `uint8`. The pixel input is a scalar.
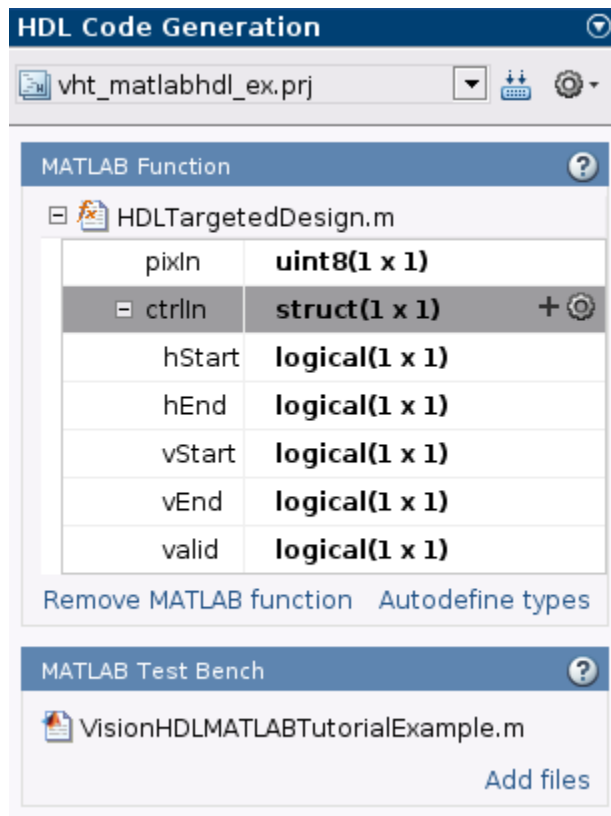
Generate HDL Code

1   Click **Workflow Advisor** to open the advisor.

2   Set **Fixed-point conversion** to `Keep original types`.

3   Click **HDL Code Generation** to view the HDL code generation options.

4   On the **Target** tab, set **Language** to `Verilog` or VHDL.

5   Also on the **Target** tab, select **Generate HDL**.

6   On the **Coding Style** tab, select **Include MATLAB source code as comments** and **Generate report** to generate a code generation report with comments and traceability links.

7   Click **Run** to generate the HDL design with reports.

8   Under **HDL Verification**, open **Verify with HDL Test Bench** and select **Generate HDL Test Bench**.

9   Click **Run** to generate the HDL test bench.

Examine the log window and click the links to view the generated code and the reports. You can run the generated test bench in your HDL simulator, or integrate the generated HDL code into a larger design you already have.

See Also

Related Examples

- "Get Started with MATLAB to HDL Workflow" (HDL Coder)
- "Generate HDL Code from MATLAB Code Using the Command Line Interface" (HDL Coder)
- "HDL Code Generation for System Objects" (HDL Coder)
- "Pixel-Streaming Design in MATLAB" on page 2-206

Modeling External Memory
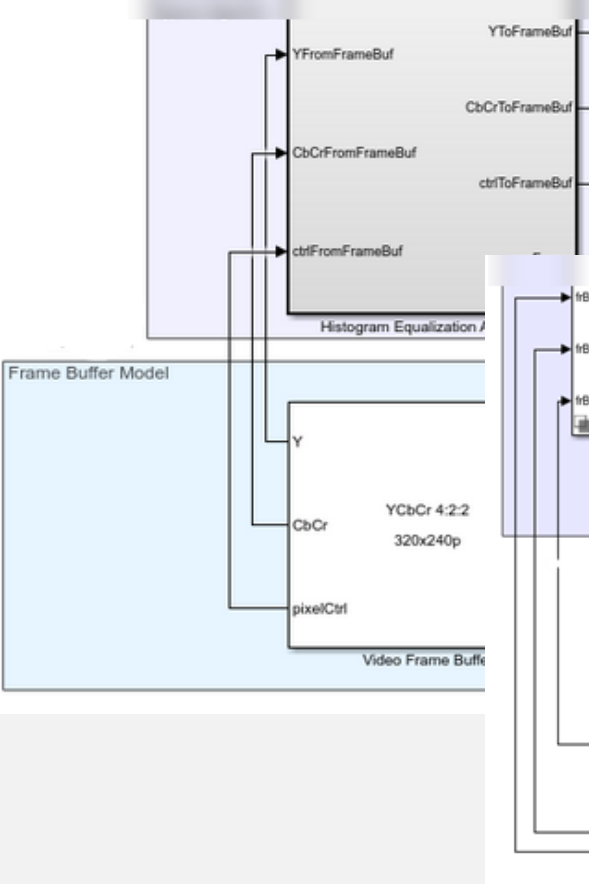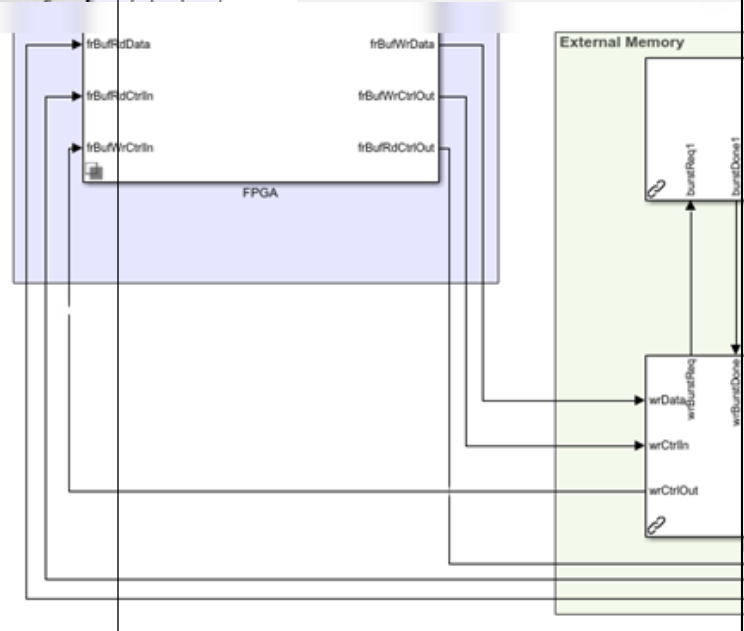
You can model external memory using features from Vision HDL Toolbox Support Package for Xilinx®
Zynq®-Based Hardware or SoC Blockset™. Both products provide models for a frame buffer or a
random access interface. They both also map your subsystem ports to physical AXI memory interfaces
when you generate HDL code and target a prototype board.

Vision HDL Toolbox Support Package for Xilinx Zynq-Based Hardware provides a simple model of the
memory interface. It does not model the timing of the interface. This level of modeling assists with
targeting a memory interface on hardware, but behavior can differ between the simulation and the
hardware. For more information, see "Model External Memory Interfaces" (Vision HDL Toolbox
Support Package for Xilinx Zynq-Based Hardware).

SoC Blockset provides library blocks to model a memory controller and multiple memory channels.
This model calculates and visualizes memory bandwidth, burst counts, and transaction latencies in
simulation. You can also model memory accesses from a processor as part of hardware-software co-
design. Use the **SoC Builder** app to generate code for FPGA and processor designs and load and run
the design on a board. You can also deploy an AXI memory interconnect monitor on your FPGA, which
can return memory transaction information for debugging and visualization in Simulink. This level of
modeling helps you verify throughput and latency requirements and enables modeling of multiple
memory consumers, including processor memory access. For more information, see "Memory" (SoC
Blockset).

Frame Buffer

Vision HDL Toolbox Support Package for Xilinx Zynq-Based Hardware	SoC Blockset
This figure shows part of the "Histogram Equalization with Zynq-Based Hardware" (Vision HDL Toolbox Support Package for Xilinx Zynq-Based Hardware) example. The Video Frame Buffer block accepts and returns the pixel streaming interface used by Vision HDL Toolbox blocks. It reads and returns an entire frame when you set the **pop** signal to 1. To use this block in your designs, copy it from the example model.	This figure shows part of the "Histogram Equalization Using Video Frame Buffer" (SoC Blockset) example. The example shows how to use the Memory Channel and Memory Controller library blocks to model a frame buffer and additional memory consumers. You can use this model to confirm that the memory interface meets the throughput and latency requirements of your design. You can measure the bandwidth and transaction latency for each memory consumer and check the measurements against the total bandwidth available from the memory. To model a frame buffer that supports the pixel streaming interface used by Vision HDL Toolbox blocks, configure the **Channel type** parameter of the Memory Channel block as AXI4 Stream Video Frame Buffer.

Random Access

Vision HDL Toolbox Support Package for Xilinx Zynq-Based Hardware	SoC Blockset
This figure shows part of the "Image Rotation with Zynq-Based Hardware" (Vision HDL Toolbox Support Package for Xilinx Zynq-Based Hardware) example. The External Memory block reads and writes to any address in the memory. In this case, rather than connecting the pixel stream to the memory interface, your custom FPGA logic must generate read and write transactions with specific addresses. To use this block in your designs, copy it from the example model.	This figure shows part of the "Random Access of External Memory" (SoC Blockset) example. This design uses a Memory Controller and two Memory Channel blocks to implement a random-access interface. In this case, rather than connecting the pixel stream to the memory interface, your custom FPGA logic must generate read and write transactions with specific addresses.

See Also

"Model External Memory Interfaces" (Vision HDL Toolbox Support Package for Xilinx Zynq-Based Hardware) | "Memory" (SoC Blockset)

Related Examples

- "Vertical Video Flipping Using External Memory" on page 3-60
- "Contrast Limited Adaptive Histogram Equalization with External Memory" on page 3-89

Deploy and Verify YOLO v2 Vehicle Detector on FPGA

This example shows how to deploy a you only look once (YOLO) v2 vehicle detector on FPGA and verify the end-to-end application using MATLAB.

The end-to-end application includes preprocessing steps, image resize and normalization, followed by a YOLO v2 vehicle detection network.

The example deploys the algorithm to a Xilinx® Zynq® Ultrascale+(TM) MPSoC ZCU102 board. Set up the board's SD card using "Guided SD Card Set Up" (Deep Learning HDL Toolbox Support Package for Xilinx FPGA and SoC Devices).

Introduction

A YOLO v2 vehicle detection application is composed of three main modules. The first module, preprocessing, accepts the input image frame and performs image resize and normalization. In the second module, the preprocessed data is consumed by the YOLO v2 vehicle detection network, which internally comprises a feature extraction network followed by a detection network. In the third module, the network output is postprocessed for identifying the strongest bounding boxes and the resulting bounding box is overlaid on the input image. In this example, as shown in the below block diagram, the first two modules are deployed on the FPGA and the postprocessing is done in MATLAB.
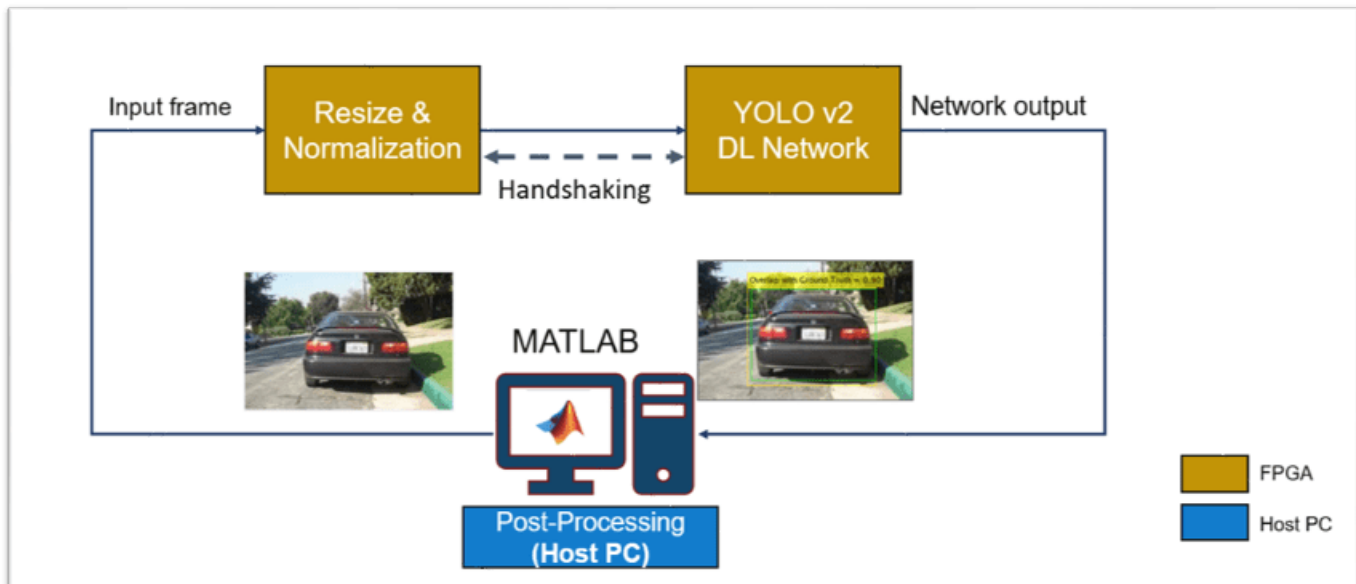
This example shows how to:

1. Configure the deep learning processor and generate IP core.
2. Model the design under test (DUT) that includes preprocessing modules (resize and normalization) and handshaking logic with the deep learning processor.
3. Generate and deploy bitstream to the FPGA.
4. Compile and deploy YOLO v2 deep learning network.
5. Verify the deployed YOLO v2 vehicle detector using MATLAB.

Configure Deep Learning Processor and Generate IP Core

The deep learning processor IP core accesses the preprocessed input from the DDR memory, performs the vehicle detection, and loads the output back into the memory. To generate a deep learning processor IP core that has the required interfaces, create a deep learning processor configuration by using the `dlhdl.ProcessorConfig (Deep Learning HDL Toolbox)` class. In the processor configuration, set the `InputRunTimeControl` and `OutputRunTimeControl` parameters. These parameters indicate the interface type for interfacing between the input and output of the deep learning processor. To learn about these parameters, see "Interface with the Deep Learning Processor IP Core" (Deep Learning HDL Toolbox). In this example, the deep learning processor uses the `register` mode for input and output runtime control.

```
hPC = dlhdl.ProcessorConfig;
hPC.InputRunTimeControl = "register";
hPC.OutputRunTimeControl = "register";
```

Specify the `TargetPlatform` property of the processor configuration object as `Generic Deep Learning Processor`. This option generates a custom generic deep learning processor IP core.

```
hPC.TargetPlatform = 'Generic Deep Learning Processor';
```

Use the `setModuleProperty` method to set the properties of the `conv` module of the deep learning processor. These properties can be tuned based on the design choice to ensure that the design fits on the FPGA. To learn more about these parameters, see `setModuleProperty (Deep Learning HDL Toolbox)`. In this example, `LRNBlockGeneration` is turned on and `SegmentationBlockGeneration` is turned off to support YOLOv2 vehicle detection network. `ConvThreadNumber` is set to 9.

```
hPC.setModuleProperty('conv','LRNBlockGeneration', 'on');
hPC.setModuleProperty('conv','SegmentationBlockGeneration', 'off');
hPC.setModuleProperty('conv','ConvThreadNumber',9);
```

This example uses the Xilinx ZCU102 board to deploy the deep learning processor. Use the `hdlsetuptoolpath` function to add the Xilinx Vivado synthesis tool path to the system path.

```
hdlsetuptoolpath('ToolName','Xilinx Vivado','ToolPath','C:\Xilinx\Vivado\2020.2\bin\vivado.bat')
```

Use the `dlhdl.buildProcessor` function with the `hPC` object to generate the deep learning IP core. It takes some time to generate the deep learning processor IP core.

```
dlhdl.buildProcessor(hPC);
```

The generated IP core contains a standard set of registers and the generated IP core report. The IP core report is generated in the same folder as ip core with the name `testbench_ip_core_report.html`.

3-15

IP core name	dlprocessor
IP core version	1.0
IP core folder	dlhdl_prj\ipcore\dlprocessor_v1_0
IP core zip file name	dlprocessor_v1_0.zip
Target platform	Generic Deep Learning Processor Xilinx
Target tool	Xilinx Vivado
Target language	VHDL
Model	testbench

`IP core name` and `IP core folder` are required in a subsequent step in 'Set Target Reference Design' task of the IP core generation workflow of the DUT. The IP core report also has the address map of the registers that are needed for handshaking with input and output of deep learning processor IP core.

Port Name	Port Type	Data Type	Target Platform Interfaces	Interface Mapping
InputNext	Inport	boolean	AXI4	x"350"
OutputNext	Inport	boolean	AXI4	x"360"
StreamingMode	Inport	boolean	AXI4	x"34C"
InputStop	Inport	boolean	AXI4	x"374"
inputStart	Inport	boolean	AXI4	x"224"
FrameCount	Inport	uint32	AXI4	x"24C"
InputValid	Outport	boolean	AXI4	x"354"
InputAddr	Outport	uint32	AXI4	x"358"
InputSize	Outport	uint32	AXI4	x"35C"
OutputValid	Outport	boolean	AXI4	x"364"
OutputAddr	Outport	uint32	AXI4	x"368"
OutputSize	Outport	uint32	AXI4	x"36C"

The registers `InputValid`, `InputAddr`, and `InputSize` contain the values of the corresponding handshaking signals that are required to write the preprocessed frame into DDR memory. The register `inputNext` is used by the DUT to pulse the inputNext signal after the data is written into memory. These register addresses are setup in the `helperSLYOLOv2PreprocessSetup.m` script. The other registers listed in the report are read/written using MATLAB. For more details on interface signals, see the Design Processing Mode Interface Signals section of "Interface with the Deep Learning Processor IP Core" (Deep Learning HDL Toolbox).

Model Design Under Test (DUT)

This section describes the design of the preprocessing modules (image resize and image normalization) and the handshaking logic in a DUT.

```
open_system('YOLOv2PreprocessTestbench');
```

YOLO v2 DUT - Preprocess with deep learning hand shake logic
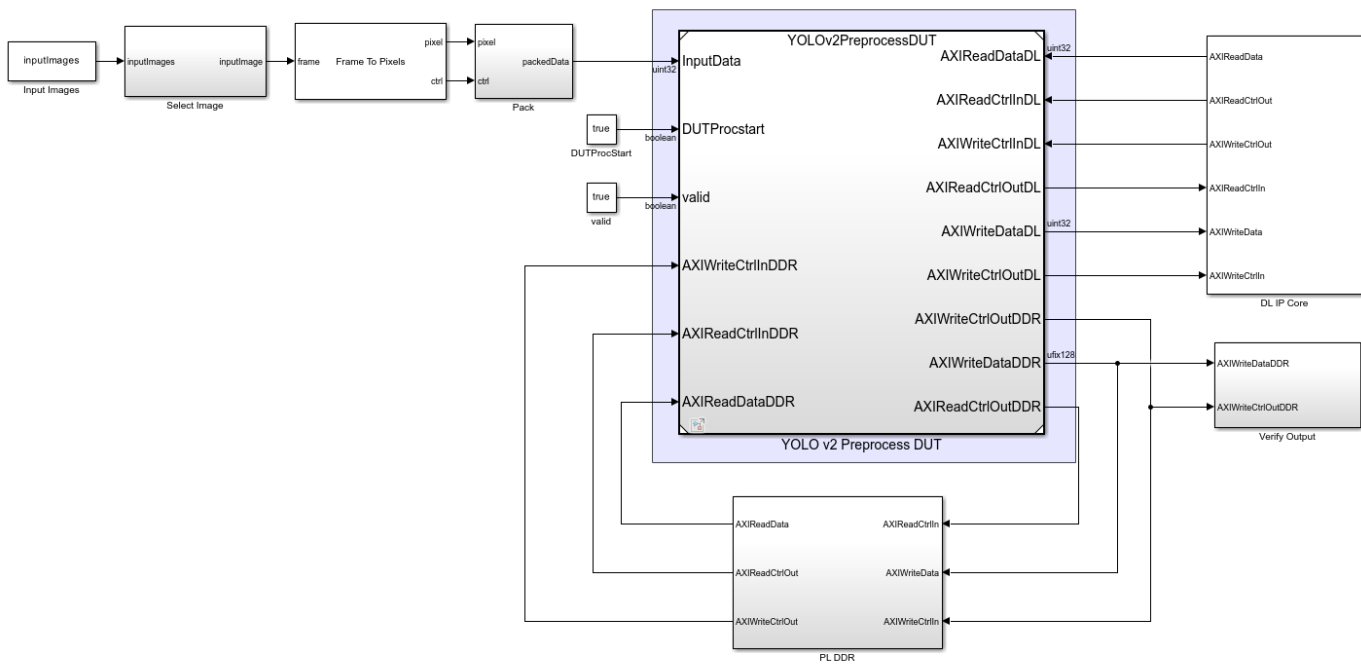
Copyright 2021-2022 The MathWorks, Inc.

The figure shows the top level view of the `YOLOv2PreprocessTestbench.slx` model. The `InitFcn` callback of the model configures the required workspace variables for the model using `helperSLYOLOv2PreprocessSetup.m` script. The `Select Image` subsystem selects the input frame from the `Input Images` block. A `Frame To Pixels` block converts the input image frame from the `Select Image` block to a pixel stream and `pixelcontrol` bus. The `Pack` subsystem concatenates the R, G, B components of the pixel stream and the five control signals of the `pixelcontrol` bus to form `uint32` data. The packed data is fed to the `YOLO v2 Preprocess DUT` for resizing and normalization. This preprocessed data is then written to the DDR using the handshaking signals from deep learning IP core. The DDR memory and the deep learning processor IP core are modeled as `PL DDR` and `DL IP core` subsystems. The model also includes a `Verify Output` subsystem which logs the signals required for the verification of the preprocessed data being written to memory using `preprocessDUTVerify.m` script.

```
open_system('YOLOv2PreprocessTestbench/YOLO v2 Preprocess DUT');
```

3-17
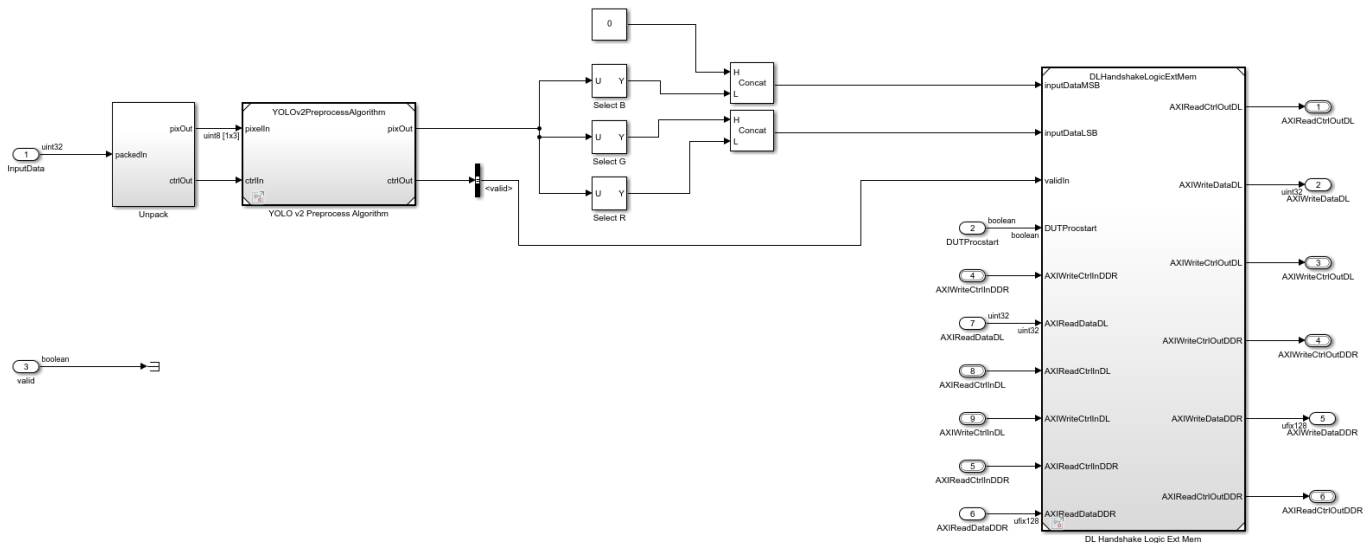
The `YOLO v2 Preprocess` DUT contains subsystems for unpacking, preprocessing (resize and normalization) and handshaking logic. The `Unpack` subsystem returns the packed input to the pixel stream and `pixelcontrol` bus. In the `YOLO v2 Preprocess Algorithm` subsystem, the input pixel stream is resized and rescaled as required by the deep learning network. This preprocessed frame is then passed to the `DL Handshake Logic Ext Mem` subsystem to be written into the PL DDR. This example models two AXI4 Master interfaces to write the preprocessed frame to the DDR memory and to read and write the registers of deep learning IP Core.

```
open_system('YOLOv2PreprocessDUT/YOLO v2 Preprocess Algorithm');
```
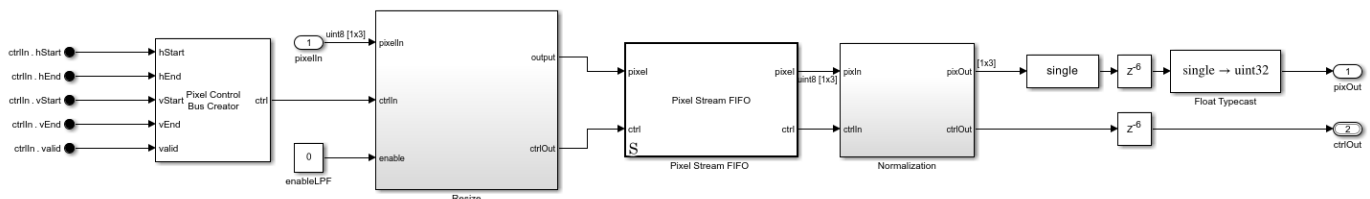

The `YOLO v2 Preprocess Algorithm` subsystem comprises of resizing, and normalization operations. The pixel stream is passed to the `Resize` subsystem for resizing to the dimensions expected by the deep learning network. The input image dimensions and the network input dimensions are setup using `helperSLYOLOv2PreprocessSetup.m` script. The resized input is passed to `Normalization` subsystem for rescaling the pixel values to [0, 1] range. The resize and normalization algorithms used in this example are described in the "Change Image Size" on page 2-175 and "Image Normalization Using External Memory" on page 3-77 examples respectively.

```
open_system('YOLOv2PreprocessDUT/DL Handshake Logic Ext Mem');
```

Deep Learning hand shake logic with external memory
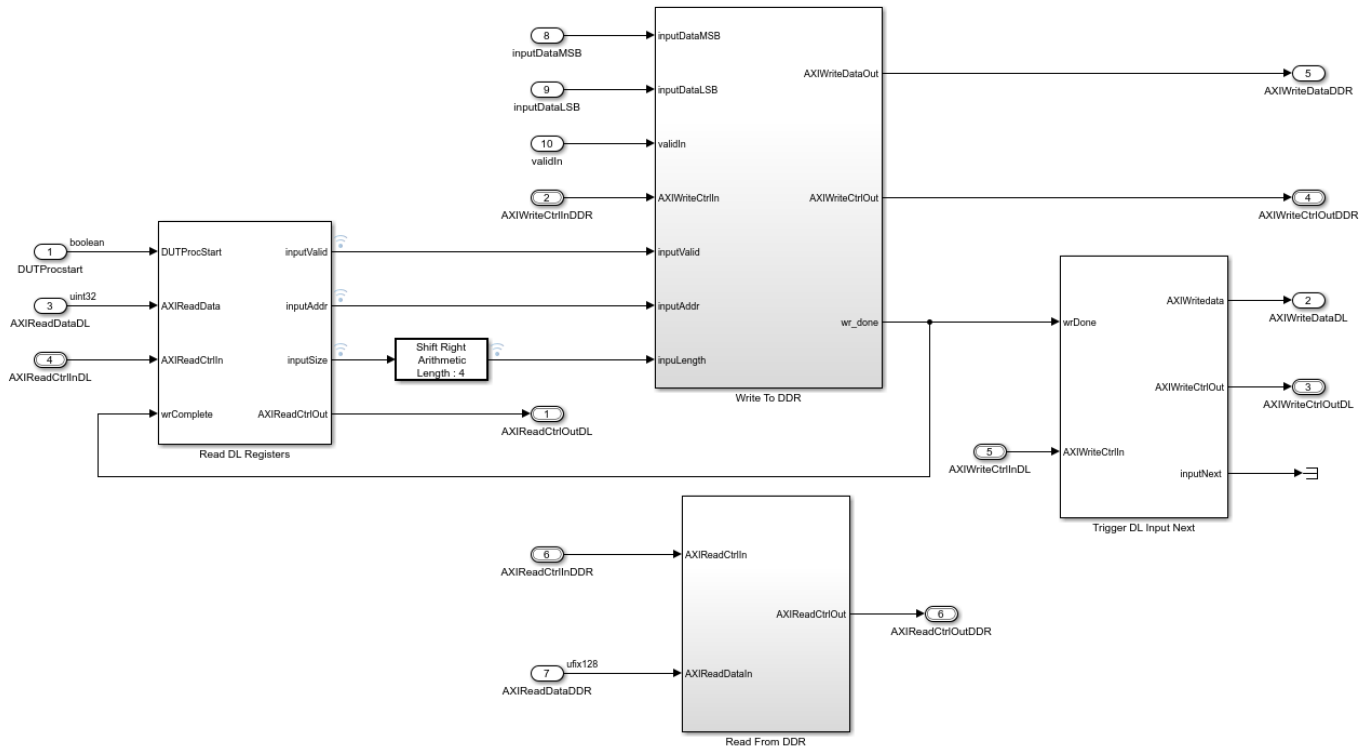
Copyright 2022 The MathWorks, Inc.

The `DL Handshake Logic Ext Mem` subsystem contains the finite state machine (FSM) logic for handshaking with DL IP and a subsystem to write the frame to DDR. The `Read DL Registers` subsystem has the FSM logic to read the handshaking signals (InputValid, InputAddr, and InputSize) from the DL IP core for multiple frames. The `Write to DDR` subsystem uses these handshaking signals to write the preprocessed frame to the memory using AXI stream protocol. The output write control bus from the DDR memory contains a signal `wr_done` which indicates that the frame write operation is done successfully. The `TriggerDLInputNext` subsystem pulses the `inputNext` signal after the preprocessed frame is written into the DDR to indicate to the DL IP core that the input data frame is available for processing.

In the next section, the IP core is generated for the `YOLO v2 Preprocess DUT` subsystem and is integrated into the reference design.

Generate and Deploy Bitstream to FPGA

This example uses the `Deep Learning with Preprocessing Interface` reference design that is provided by the Vision HDL Toolbox™ Support Package for Xilinx® Zynq®-Based Hardware.

```
pathToRefDesign = fullfile(...
    matlabshared.supportpkg.getSupportPackageRoot,...
    "toolbox","shared","supportpackages","visionzynq","target",....
    "+visionzynq", "+ZCU102", "plugin_rd.m");
if (~exist(pathToRefDesign, 'file'))
    error(['This example requires you to download and install '...
```

```
        'Vision HDL Toolbox Support Package for Xilinx Zynq-Based Hardware']);
end
```

The reference design contains the ADI AXI DMA Controller to move the data from processor to FPGA fabric. The data is sent from the ARM processing system, through the DMA controller and AXI4-Stream interface, to the generated DUT Preprocessing IP core. The DUT contains two AXI Master interfaces. One AXI interface is connected to the Deep Learning Processor IP core and the other is connected to the DDR memory interface generator (MIG) IP.
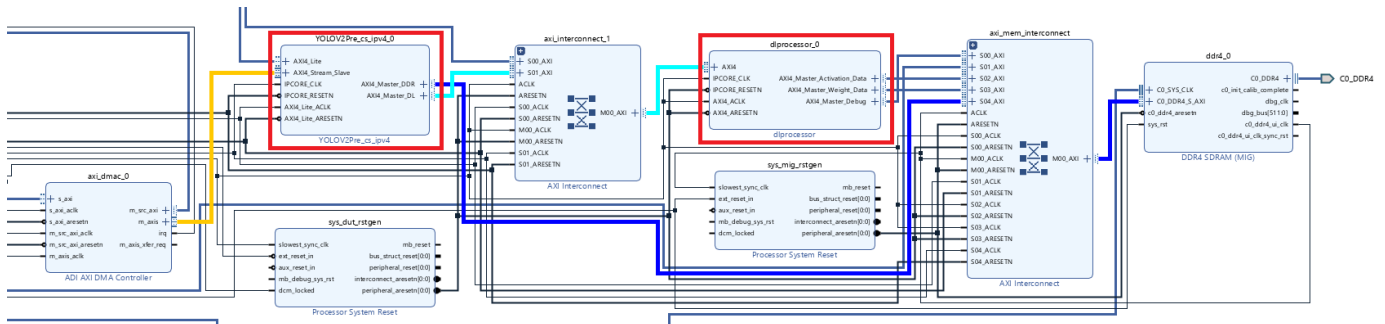
Start the targeting workflow by right clicking the `YOLO v2 Preprocess DUT` subsystem and selecting HDL Code > HDL Workflow Advisor.

- In step 1.1, select IP Core Generation workflow and the platform 'Xilinx Zynq Ultrascale+ MPSoC ZCU102 Evaluation Kit'.
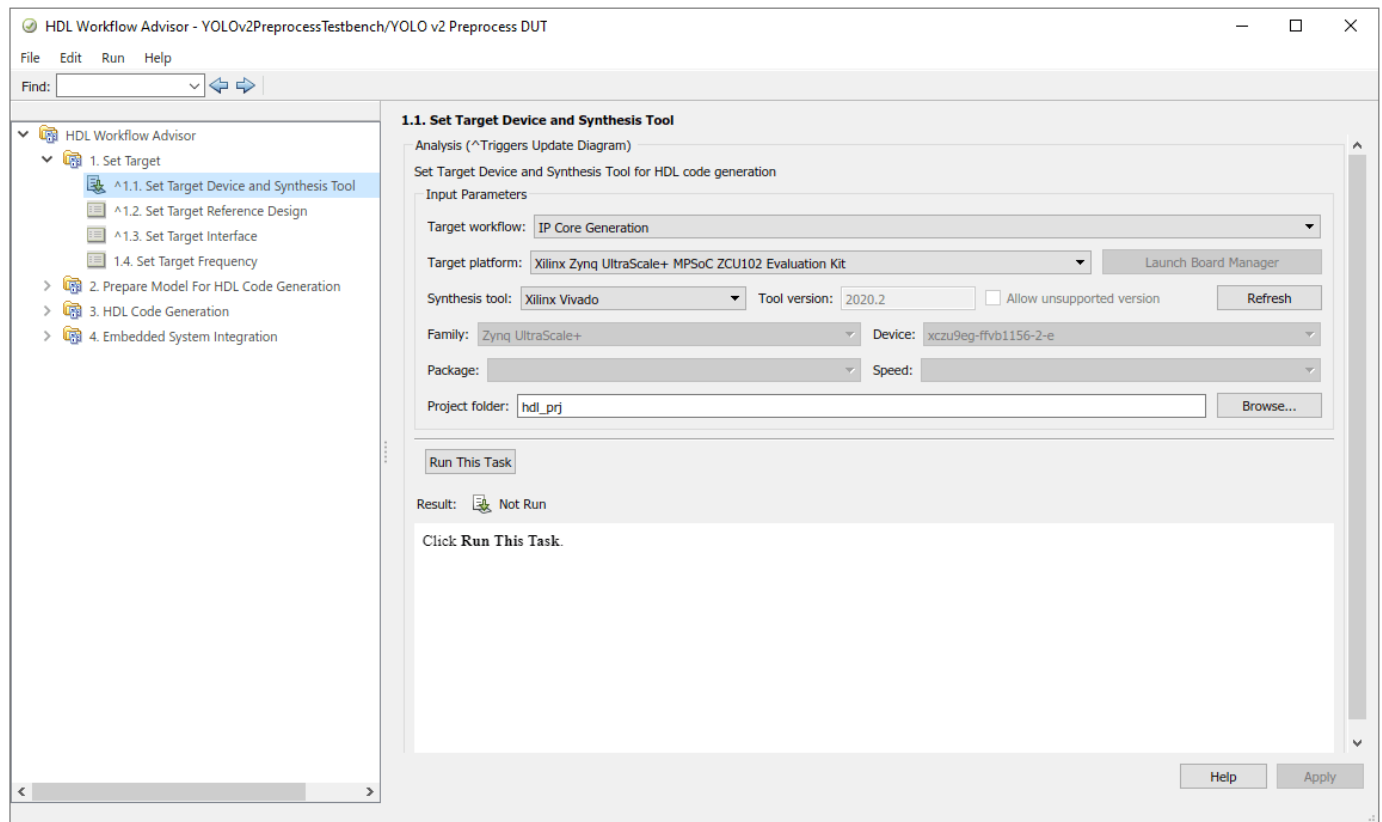
- In step 1.2, the reference design is set to "Deep Learning with Preprocessing Interface". The DL Processor IP name and the DL Processor IP location specify the name and location of the generated deep learning processor IP core, and are obtained from the IP core report.
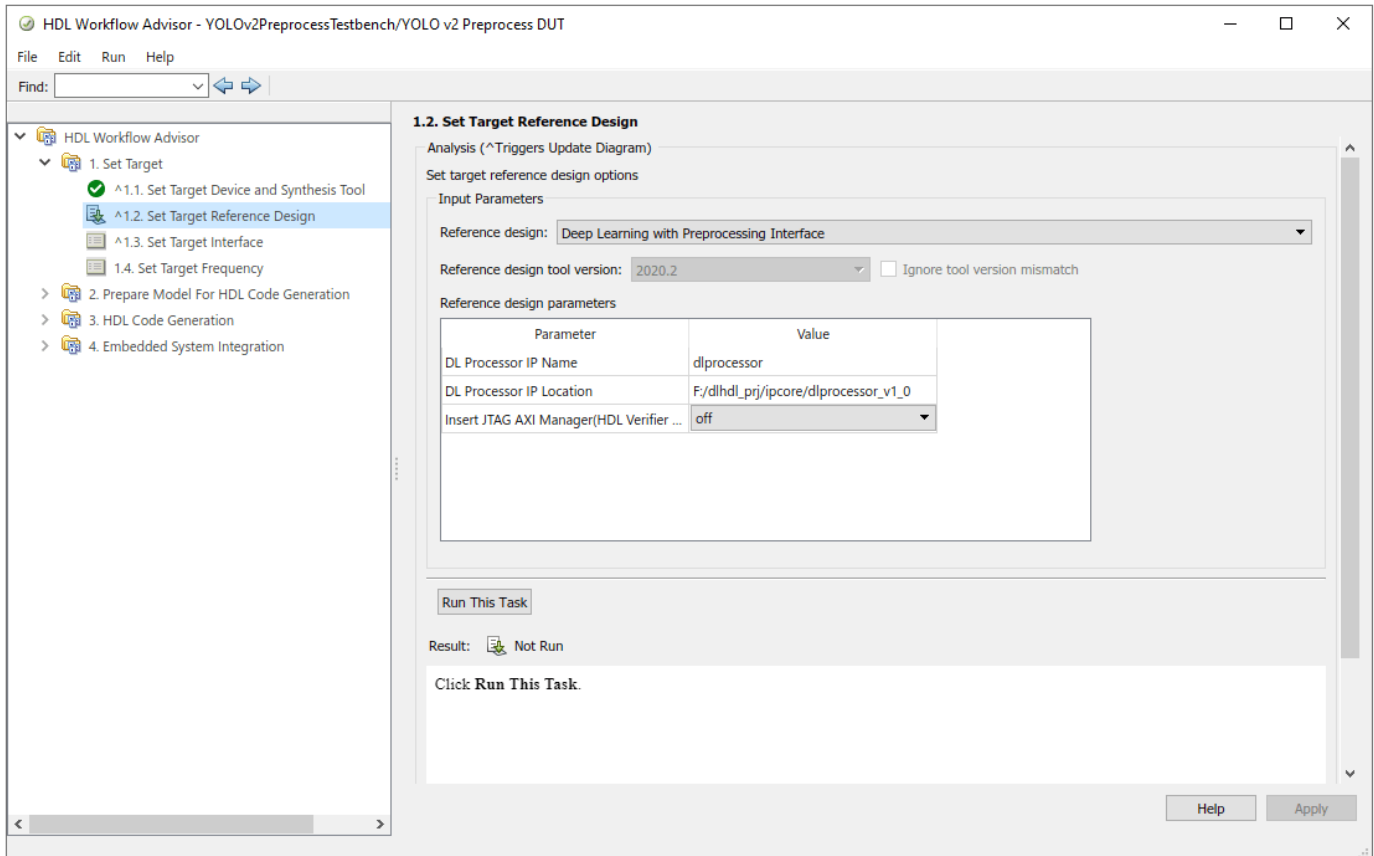
- In step 1.3, map the target platform interfaces to the input and output ports of the DUT.

Port Name	Port Type	Data Type	Target Platform Interfaces	Interface Mapping	Interface Options
inputData	Inport	uint32	AXI4-Stream Slave ▾	Data ▾	
DUTProcstart	Inport	boolean	AXI4-Lite ▾	x"100"	Options...
valid	Inport	boolean	AXI4-Stream Slave ▾	Valid ▾	
AXIWriteCtrlInDDR	Inport	bus	AXI4 Master DDR Write ▾	Write Slave to Master B ▾	
AXIReadCtrlInDDR	Inport	bus	AXI4 Master DDR Read ▾	Read Slave to Master Bu ▾	
AXIReadDataDDR	Inport	ufix128	AXI4 Master DDR Read ▾	Data ▾	
AXIReadDataDL	Inport	uint32	AXI4 Master DL Read ▾	Data ▾	
AXIReadCtrlInDL	Inport	bus	AXI4 Master DL Read ▾	Read Slave to Master Bu ▾	
AXIWriteCtrlInDL	Inport	bus	AXI4 Master DL Write ▾	Write Slave to Master B ▾	
AXIReadCtrlOutDL	Outport	bus	AXI4 Master DL Read ▾	Read Master to Slave Bu ▾	
AXIWriteDataDL	Outport	uint32	AXI4 Master DL Write ▾	Data ▾	
AXIWriteCtrlOutDL	Outport	bus	AXI4 Master DL Write ▾	Write Master to Slave B ▾	
AXIWriteCtrlOutDDR	Outport	bus	AXI4 Master DDR Write ▾	Write Master to Slave B ▾	
AXIWriteDataDDR	Outport	ufix128	AXI4 Master DDR Write ▾	Data ▾	
AXIReadCtrlOutDDR	Outport	bus	AXI4 Master DDR Read ▾	Read Master to Slave Bu ▾	

- **AXI4-Stream Slave interface**: The `inputData` and `valid` ports of the DUT are mapped to the data and valid ports of the AXI4-Stream Slave interface respectively.
- **AXI4-Lite Interface**: The `DUTProcstart` register is mapped to the AXI4-Lite register. When this register is written, it triggers the process of input handshaking logic. Choosing the AXI4-Lite interface directs HDL Coder to generate a memory-mapped register in the FPGA fabric. You can access this register from software running on the ARM processor.
- **AXI4 Master DDR interface**: The `AXIWriteCtrlInDDR`, `AXIReadCtrlInDDR`, `AXIReadDataDDR`, `AXIWriteCtrlOutDDR`, `AXIWriteDataDDR` and `AXIReadCtrlOutDDR` ports of DUT are mapped to AXI4 Master DDR interface. The **Read Channel** of the AXI4 Master DDR interface is mapped to the AXI4 Master DDR Read interface, and the **Write Channel** of the AXI4 Master DDR interface is mapped to the AXI4 Master DDR Write interface. This interface is used for the data transfer between the Preprocess DUT and the PL DDR. Using the Write Channel of this interface, the preprocessed data is written to the PL DDR which can then be accessed by the Deep Learning Processor IP.
- **AXI4 Master DL interface**: The `AXIReadDataDL`, `AXIReadCtrlInDL`, `AXIWriteCtrlInDL`, `AXIReadCtrlOutDL`, `AXIWriteDataDL` and `AXIWriteCtrlOutDL` ports of DUT are mapped to AXI4 Master DL interface. The **Read Channel** of the AXI4 Master DL interface is mapped to the AXI4 Master DL Read interface, and the **Write Channel** of the AXI4 Master DL interface is mapped to the AXI4 Master DL Write interface. This interface is used for the communication between Preprocess DUT and the Deep Learning Processor IP. In this example, this interface is used for implementing input handshaking logic with Deep Learning Processor IP.
- Step 2 prepares the design for hdl code generation.
- Step 3 generates HDL code for the IP core.
- Step 4.1 integrates the newly generated IP core into the reference design.
- In step 4.2, the host interface script and Zynq software interface model is created. Since this example uses the interface script, and not the model, uncheck **Generate Simulink software**

interface model. The host interface script, `gs_YOLOv2PreprocessTestbench_interface`, generated in this step is parameterized and provided as `setupPreprocessIPInterfaces.m` function as part of this example.
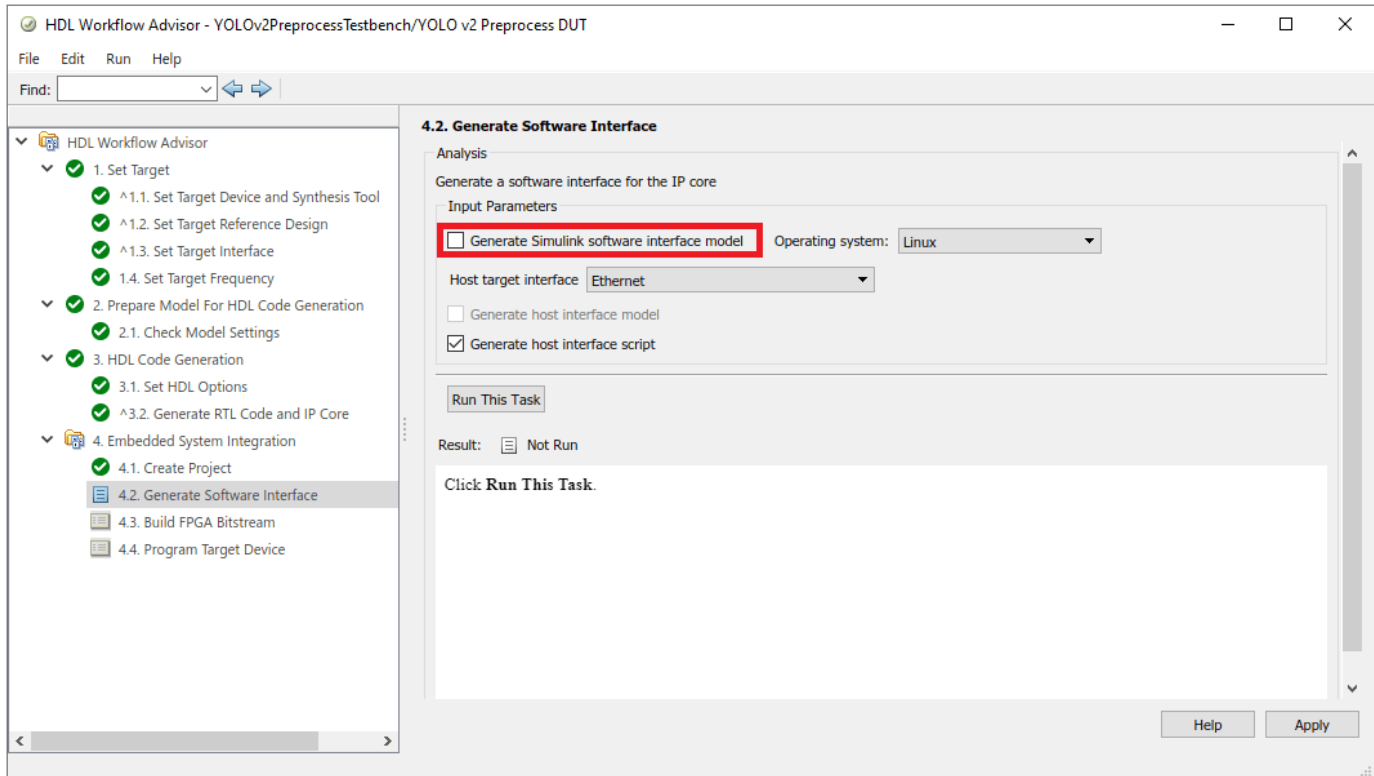
- Step 4.3 generates the bitstream. The bit file is named `block_design_wrapper.bit` and located at hdl_prj\vivado_ip_prj\vivado_prj.runs\impl_1. This bitstream is downloaded to FPGA in the next section.

Compile and Deploy Yolo v2 Deep Learning Network

Now that the bitstream is generated for the IP core of the DUT integrated with the reference design that contains the DL IP core, you can deploy the end to end deep learning application onto an FPGA.

Create a target object to connect your target device to the host computer. Use the installed Xilinx Vivado Design Suite over an Ethernet connection to program the device.

```
hTarget = dlhdl.Target('Xilinx','Interface','Ethernet','IpAddr','192.168.1.101');
```

Load the pretrained YOLO v2 object detection network.

```
vehicleDetector = load('yolov2VehicleDetector.mat');
detector = vehicleDetector.detector;
net = detector.Network;
```

Update the bitstream build information in the MAT file generated during the IP core generation. The name of the MAT file is `dlprocessor.mat` and is located in cwd\dlhdl_prj\, where cwd is your current working folder. Copy the file to the present working folder. This MAT file generated using the target platfom `Generic Deep Learning Processor` does not contain the Board/Vendor

information. Use `updateBitstreamBuildInfo.m` function to update the Board/Vendor information and generate a new MAT file with the same name as generated bitstream.

```
bitstreamName = 'block_design_wrapper';
updateBitstreamBuildInfo('dlprocessor.mat',[bitstreamName,'.mat']);
```

Create a deep learning HDL workflow object using the `dlhdl.Workflow` class.

```
hW = dlhdl.Workflow('Network',net,'Bitstream',[bitstreamName,'.bit'],'Target',hTarget);
```

Compile the network, `net` using the `dlhdl.Workflow` object.

```
frameBufferCount = 3;
compile(hW, 'InputFrameNumberLimit', frameBufferCount);
```

Create a Xilinx processor hardware object and connect to the processor on-board the Xilinx SoC board.

```
hSOC = xilinxsoc('192.168.1.101', 'root', 'root');
```

Call the `xilinxsoc` object function `ProgramFPGA` to program the FPGA and set the device tree to use the processor on the SoC board.

```
programFPGA(hSOC, [bitstreamName,'.bit'], 'devicetree_vision_dlhdl.dtb');
```

Run the `deploy` function of the `dlhdl.Workflow` object to download the network weights and biases on the Zynq UltraScale+ MPSoC ZCU102 board.

```
deploy(hW, 'ProgramBitStream', false);
```

Clear the DLHDL workflow object and hardware target.

```
clear hW;
clear hTarget;
```

Verify Deployed YOLO v2 Vehicle Detector Using MATLAB

The function `YOLOv2DeployAndVerifyDetector` takes `hSOC` object as input and performs vehicle detection using the YOLO v2 network deployed on FPGA and verifies the end-to-end application using MATLAB.

```
YOLOv2DeployAndVerifyDetector(hSOC);
```

This flowchart shows the operations performed in the function.

```
                              ( Start )
                                 |
                                 v
  1. unzip([datasetLocation,'\vehicleDatasetImages.zip']);         Load the vehicle data set
  2. data= load([datasetLocation,'\vehicleDatasetGroundTruth.mat']);
  3. vehicleDataset = data.vehicleDataset;
  4. inputDataTbl = vehicleDataset(153:259,:);
                                 |
                                 v
  5. hFPGA = fpga("Xilinx");                                        Setup deep learning and
  6. addrMap = setupDLIPInterfaces(hFPGA,bitstream,readFrameLength); preprocessing interfaces.
  7. ddrbaseAddr = dec2hex(addrMap('ddrbase'));
  8. setupPreprocessIPInterfaces(hFPGA,writeFrameLength);
                                 |
                                 v
  8.  writePort(hFPGA,"StreamingMode",1);                           Configure deep learning IP core in
  10. writePort(hFPGA,"FrameCount",0);                              continuous streaming mode.
  11. Pulse inputStart signal
                                 |
                                 v
  12. for imIndex = 1:numFrames:                                    Pack control and data into 32-bit
  13. inputImage = imread(inputDataTbl.imageFilename{imIndex});     data and send the packed input to
  14. packedIn = packPixelAndControlBus(inputImage,frm2pix);        the YOLO v2 Preprocess DUT for
  15. writePort(hFPGA,"InputData",packedIn);                        each frame.
                                 |
                                 v
  16. outData = readDataFromPLDDR(hFPGA, ddrbaseAddr);              Read the output data from memory.
  17. Pulse outputNext signal                                       Perform the YOLO v2 transform
  18. [bboxes,~] = yolov2TransformLayerAndPostProcess(outData,...   layer and post-processing.
        inputImage,networkInputSize,networkOutputSize,anchorBoxes);
                                 |
                                 v
  19. bboxesGT = inputDataTbl.vehicle{imIndex};                     Verify the output obtained after
  20. overlappingRatio = bboxOverlapRatio(bboxes, bboxesGT);        post-processing using the ground truth
  21. Pulse inputStop signal                                        for each frame and pulse inputStop
                                                                    signal
```

This section describes the steps in the flowchart in detail.

Load the vehicle data set

```
datasetLocation = [matlabroot, filesep, 'examples', filesep, 'deeplearning_shared', filesep, 'da
unzip([datasetLocation, filesep, 'vehicleDatasetImages.zip']);
data = load([datasetLocation, filesep, 'vehicleDatasetGroundTruth.mat']);
vehicleDataset = data.vehicleDataset;
```

The vehicle data is stored in a two-column table, where the first column contains the image file paths and the second column contains the vehicle bounding boxes. Add the fullpath to the local vehicle data folder.

```
vehicleDataset.imageFilename = fullfile(pwd,vehicleDataset.imageFilename);
```

Select images from the vehicle dataset. Each image present in `inputDataTbl` has 224 rows and 340 columns.

```
inputDataTbl = vehicleDataset(153:259,:);
```

Setup deep learning and preprocessing interfaces

Connect to the FPGA on-board the SoC board by using the fpga function. Use the processor hardware object `hSOC` as an input to the `fpga` function.

```
hFPGA = fpga(hSOC);
```

Get network input and output size. The `networkOutputSize` is the output size of `yolov2ClassConv` obtained from `analyzeNetwork(net)`.

```
networkInputSize = net.Layers(1, 1).InputSize;
networkOutputSize = [16,16,24];
```

The deep learning processor writes the `yolov2ClassConv` layer output to the external memory in a specified data format. This data format depends on the chosen `ConvThreadNumber` of the deep learning processor. `readLengthFromDLIP` contains the output data size. For more information, see "External Memory Data Format" (Deep Learning HDL Toolbox)

```
readLengthFromDLIP = (networkOutputSize(1)*networkOutputSize(2)*networkOutputSize(3)*4)/3;
```

Setup the deep learning IP interfaces using `setupDLIPInterfaces.m` function. This function uses `BitstreamManager` class to obtain the address map of the deep learning IP core registers.

```
addrMap = setupDLIPInterfaces(hFPGA, [bitstreamName,'.bit'], readLengthFromDLIP);
ddrbaseAddr = dec2hex(addrMap('ddrbase'));
```

Get image dimensions and create `visionhdl.FrameToPixels` System object™

```
frm = imread(inputDataTbl.imageFilename{1});
```

```
frmActivePixels = size(frm,2);
frmActiveLines = size(frm,1);
```

```
frm2pix = visionhdl.FrameToPixels(...
      'NumComponents',size(frm,3),...
      'VideoFormat','custom',...
      'ActivePixelsPerLine',frmActivePixels,...
      'ActiveVideoLines',frmActiveLines,...
      'TotalPixelsPerLine',frmActivePixels+10,...
      'TotalVideoLines',frmActiveLines+10,...
      'StartingActiveLine',6,...
      'FrontPorch',5);
```

Setup the preprocess IP interfaces using `setupPreprocessIPInterfaces.m` function.

```
inputFrameLength = frm2pix.TotalPixelsPerLine * frm2pix.TotalVideoLines;
setupPreprocessIPInterfaces(hFPGA, inputFrameLength);
```

Configure deep learning IP core

Set data processing mode to continuous streaming mode by setting `StreamingMode` register to `true` and `FrameCount` register to `0`.

```
writePort(hFPGA, "StreamingMode", 1);
writePort(hFPGA, "FrameCount", 0);
```

Pulse the `inputStart` signal to indicate to the deep learning IP core to start processing the data.

```
writePort(hFPGA, "inputStart", 0);
writePort(hFPGA, "inputStart", 1);
writePort(hFPGA, "inputStart", 0);
```

Assert `DUTProcStart` to signal preprocess DUT to start writing the preprocessed data to the DDR.

```
writePort(hFPGA, "DUTProcStart", 1);
```

Send input video frame to YOLO v2 Preprocess DUT

Use `packPixelAndControlBus.m` function to pack the pixel and control bus data. The `frm2pix` object converts the input video frame to a pixel stream and `pixelcontrol` bus. Then, the R, G, B components of the pixel data and the `hStart`, `hEnd`, `vStart`, `vEnd`, and `valid` signals of the `pixelcontrol` bus are packed to generate 32 bit data, as shown.

This packed input is fed to the `YOLO v2 Preprocess DUT` using the `writePort` function of `fpga` object. The input is preprocessed and written to the memory by the DUT. The deep learning IP core reads the data from memory, performs the vehicle detection, and writes the output back to the memory.

```
writePort(hFPGA, "InputData", inputImagePacked);
```

Read output data and perform postprocessing

The deep learning IP core returns handshaking signals indicating address, size, and validity of the output. When the `outputValid` signal becomes `true`, the script reads the processed output data frame using the `outputAddr` and `outputSize` signals. The `readDataFromPLDDR.m` function reads the output data using the `readPort` function of `fpga` object.

```
outputValid = readPort(hFPGA, "OutputValid");
while(outputValid~=1)
    pause(0.1);
    outputValid = readPort(hFPGA, "OutputValid");
end
```

```
outData = readDataFromPLDDR(hFPGA, ddrbaseAddr);
```

After reading the output data from DDR, pulse the OutputNext signal by using the `hFPGA` object

```
writePort(hFPGA, "OutputNext", 0);
writePort(hFPGA, "OutputNext", 1);
writePort(hFPGA, "OutputNext", 0);
```

The `yolov2TransformLayerAndPostProcess.m` function performs the transform layer processing and postprocessing on the `outData`, and returns the bounding boxes.

```
anchorBoxes = detector.AnchorBoxes;
[bboxes, scores] = yolov2TransformLayerAndPostProcess(outData, inputImage, networkInputSize, netw
```

Verify postprocessed output

The bounding boxes obtained from the post-processing and the ground truth are overlaid on the input image along with the overlap ratio.

```
bboxesGT = inputDataTbl.vehicle{imIndex};
```

```
overlapRatio = bboxOverlapRatio(bboxes, bboxesGT);

bbOverlap = sprintf("Overlap with Ground Truth = %0.2f", overlapRatio);
outputImage = insertObjectAnnotation(inputImage,'rectangle',bboxes,bbOverlap);
outputImage = insertObjectAnnotation(outputImage,'rectangle',bboxesGT, '', 'Color', 'green');
imshow(outputImage);
```
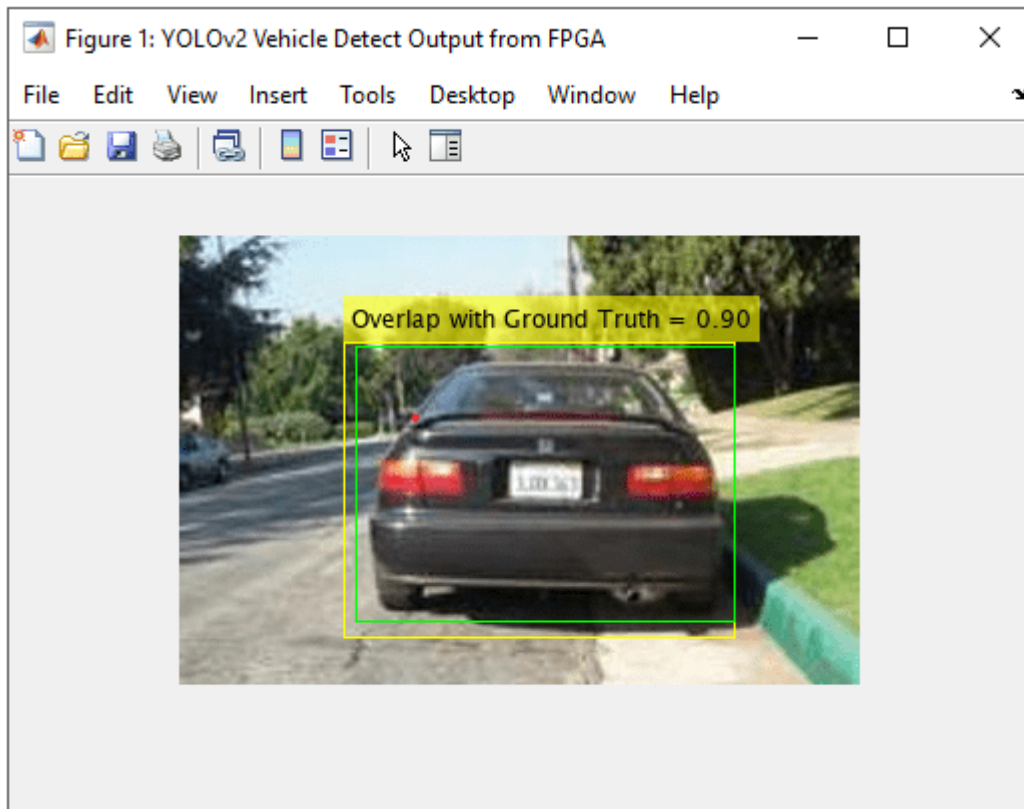

Pulse inputStop signal

After processing all the frames, pulse the inputStop signal by using the `hFPGA` object.

```
writePort(hFPGA, "InputStop", 0);
writePort(hFPGA, "InputStop", 1);
writePort(hFPGA, "InputStop", 0);
```

Conclusion

This example deployed the YOLO v2 vehicle detector application comprising of preprocessing steps (image resize and normalization) and handshaking logic on FPGA, performed vehicle detection, and verified the results using MATLAB.

For information about debugging the design deployed on the FPGA, see the "Debug YOLO v2 Vehicle Detector on FPGA" on page 3-30 example. This example shows how to use FPGA data capture and

AXI manager features of the HDL Verifier™ product to capture the required data for debugging from the FPGA.

See Also

Related Examples
- "Integrate YOLO v2 Vehicle Detector System on SoC" on page 3-41

More About
- "Target Deep Learning Processor and Image Preprocessing to FPGA" (Vision HDL Toolbox Support Package for Xilinx Zynq-Based Hardware)

Debug YOLO v2 Vehicle Detector on FPGA

This example shows how to debug hardware by visualizing signals from a vehicle detector design deployed on the Xilinx® Zynq® UltraScale+™ MPSoC ZCU102 board. You use FPGA data capture and AXI manager features of HDL Verifier™ support package for Xilinx FPGA Boards software to set triggers and capture the signals of interest. The "Deploy and Verify YOLO v2 Vehicle Detector on FPGA" on page 3-14 example shows how to deploy a vehicle detector design on an FPGA. In this example, you integrate FPGA data capture and AXI manager features into this design to debug and visualize its functionality.

Introduction

Debugging designs, especially those deployed to the FPGA, can be a difficult task without a proper set of tools. FPGA data capture and AXI manager offer many capabilities to easily debug designs deployed to an FPGA. In this example, you focus on the **Preprocessing module** of the design. You analyze several scenarios where proper debugging is required to ensure the application behaves correctly. The scenarios are:

- Handshaking between the Preprocessing DUT and deep learning (DL) IP core. This scenario shows how to use FPGA data capture and AXI manager features to visualize the handshaking events between the Preprocessing DUT and the DL IP in the Logic Analyzer (DSP System Toolbox). You use FPGA data capture to tap the handshaking signals between the Preprocessing DUT and the DL IP from the FPGA.

- Functionality of the `Resize` Subsystem. This scenario shows how to add debug hooks to the model and use them for debugging and verification.

- Handshaking between the Preprocessing DUT and DDR memory. This scenario shows how to visualize the handshaking events between the Preprocessing DUT and the DDR memory in the Logic Analyzer. You use FPGA data capture to tap the handshaking signals between the Preprocessing DUT and the DDR memory from the FPGA.

Add Debug Hooks and Test Points in Model

To capture signal data using FPGA data capture, configure the signal as a test point. For more information, see "Configure Signals as Test Points" (Simulink). Configure all the signals described in this section as test points. Use the Bus Selector (Simulink) block to extract signals from a bus and then add test points. To calculate the valid pixel flow through the `Resize` subsystem, add debugging logic using counters within the `YOLOv2PreprocessAlgorithm` model. Use the `helperConfigAndAddTestPoints` function to automate the process of adding the counters and test points to the `YOLOv2PreprocessAlgorithm` and `DLHandshakeLogicExtMem` models. The `helperConfigAndAddTestPoints` function creates the four models, which are `YOLOv2PreprocessTbDebug`, `YOLOv2PreprocessDUTDebug`, `YOLOv2PreprocessAlgoDebug`, and `DLHandshakeLogicDebug`. These four models contain all the required testpoints and debug hooks.

This figure shows the signals that are configured as test points in the `YOLOv2PreprocessAlgoDebug` model.
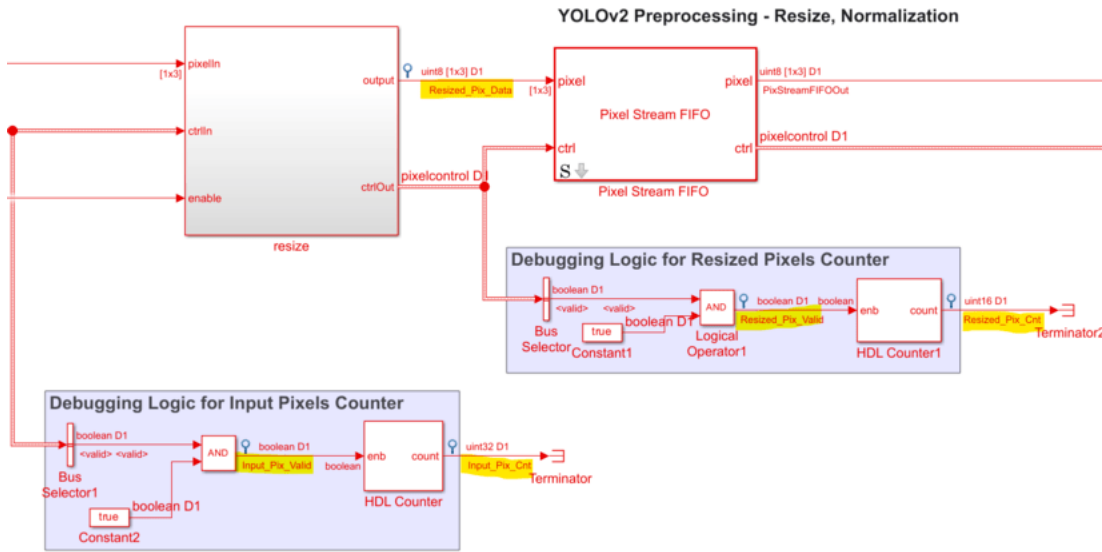
This figure shows the signals that are configured as test points in the `DLHandshakeLogicDebug` model.
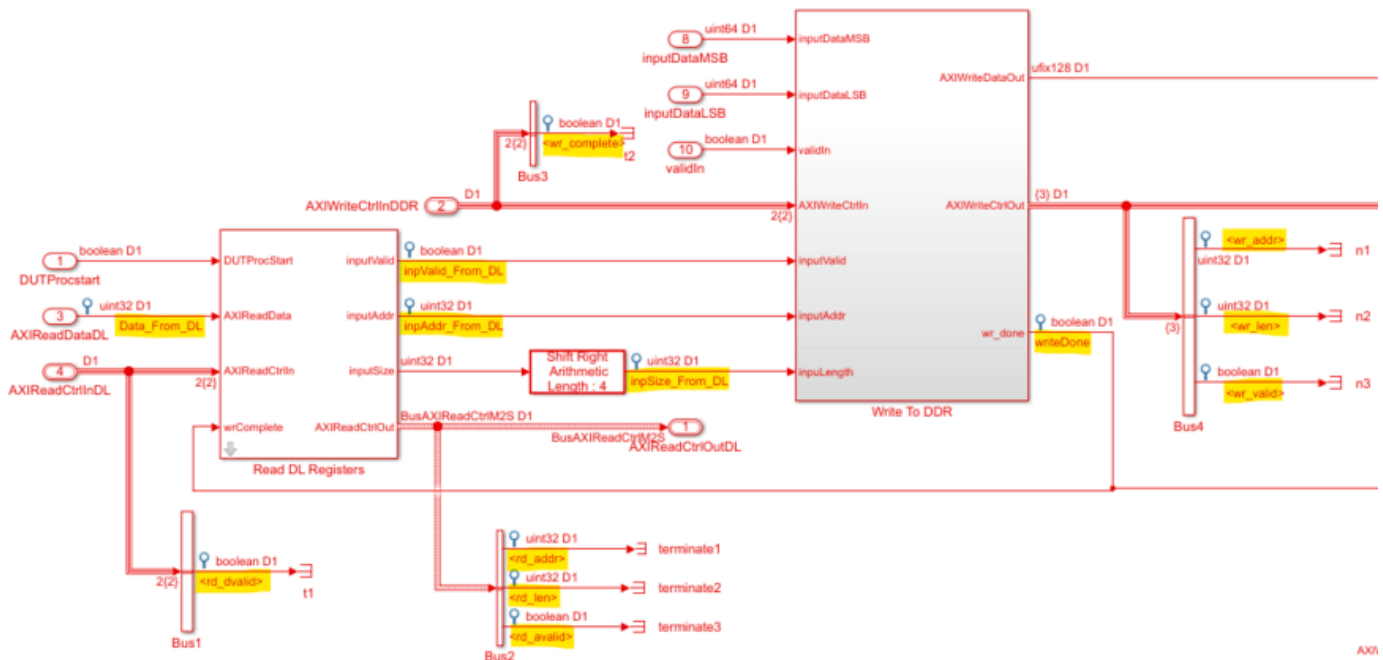
Use the `Simulink.BlockDiagram.arrangeSystem` (Simulink) function to improve the layout of the model.

Integrate FPGA Data Capture and AXI Manager in HDL Workflow Advisor

To generate IP core files for a DL processor, follow the steps in the Configure Deep Learning Processor and Generate IP Core section of the "Deploy and Verify YOLO v2 Vehicle Detector on FPGA" on page 3-14 example. Use the `helperUpdateHDLWorkflowAdvisor` function to automate the

process of configuring the HDL workflow advisor settings and generate the bitstream. You must provide the complete path to the DL IP core files. Set the buffer size for FPGA data capture IP to 16384 and the maximum sequence depth to 7.

```
pathToDLIPFiles = 'F:\dlhdl_prj\ipcore\dlprocessor_v1_0';
modelWithTestPoints = {'YOLOv2PreprocessTbDebug','YOLOv2PreprocessDUTDebug','YOLOv2PreprocessAlgo
helperUpdateHDLWorkflowAdvisor(pathToDLIPFiles,modelWithTestPoints,'16384','7')
```

Follow these steps to perform this task manually.

1   Start the targeting workflow by right-clicking the `YOLO v2 Preprocess DUT Subsystem` subsystem in the `YOLOv2PreprocessTbDebug` model and selecting **HDL Code > HDL Workflow Advisor**.

2   In step 1.1, select `IP Core Generation` and set **Target platform** to `Xilinx Zynq Ultrascale+ MPSoC ZCU102 Evaluation Kit`.

3   In step 1.2, set **Reference design** to `Deep Learning with Preprocessing Interface`. The **DL Processor IP name** and the **DL Processor IP location** fields specify the name and location of the generated deep learning processor IP core, respectively. These details are fetched from the IP core report. Set **Insert AXI manager** to `JTAG`.

4   In step 1.3, enable the **Enable HDL DUT output port generation for test points** setting to update the interface table with all the test points as output ports for the generated DUT. Map the target platform interfaces to the input and output ports of the DUT. For the required interface mapping, see step 1.3 in Generate and Deploy Bitstream to FPGA section of the "Deploy and Verify YOLO v2 Vehicle Detector on FPGA" on page 3-14 example. This table shows the interface mapping for test points. To capture and visualize the trigger signals in the Logic Analyzer, map the trigger signals to `Trigger and Data` instead of `Trigger`. For more information, see "Use As" (HDL Verifier).

Port Name	Port Type	Data Type	Target Platform Interfaces	Interface Mapping	Interface Options
Data_from_DL	Test point	uint32	FPGA Data Capture	Data	
inpSize_from_DL	Test point	uint32	FPGA Data Capture	Data	
<wr_complete>	Test point	boolean	FPGA Data Capture	Trigger and Data	
<rd_dvalid>	Test point	boolean	FPGA Data Capture	Trigger and Data	
<rd_addr>	Test point	uint32	FPGA Data Capture	Data	
<rd_len>	Test point	uint32	FPGA Data Capture	Data	
<rd_avalid>	Test point	boolean	FPGA Data Capture	Trigger and Data	
<wr_addr>	Test point	uint32	FPGA Data Capture	Data	
<wr_len>	Test point	uint32	FPGA Data Capture	Data	
<wr_valid>	Test point	boolean	FPGA Data Capture	Trigger and Data	
inpValid_from_DL	Test point	boolean	FPGA Data Capture	Trigger and Data	
inpAddr_from_DL	Test point	uint32	FPGA Data Capture	Data	
writeDone	Test point	boolean	FPGA Data Capture	Trigger and Data	
Input_Pix_Cnt	Test point	uint32	FPGA Data Capture	Data	
Resized_Pix_Cnt	Test point	uint16	FPGA Data Capture	Data	
Resized_Pix_Valid	Test point	boolean	FPGA Data Capture	Trigger and Data	
Inp_Pix_Valid	Test point	boolean	FPGA Data Capture	Trigger and Data	
Resized_Pix_Data	Test point	uint8 (3)	FPGA Data Capture	Data	

- Perform steps 1.4 to 3.1 as shown in the Generate and Deploy Bitstream to FPGA section of the "Deploy and Verify YOLO v2 Vehicle Detector on FPGA" on page 3-14 example.

- In step 3.2, set **FPGA data capture buffer size** to 16384 and **FPGA data capture maximum sequence depth** to 7. Select **Include capture condition logic in FPGA data capture** to enable the capture control logic option in the generated FPGA data capture component.
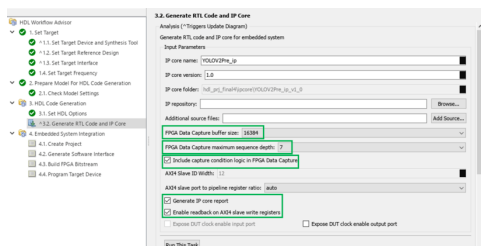
- In step 4.3, generate the bitstream. The HDL Workflow Advisor generates the `block_design_wrapper.bit` bitstream file in the `hdl_prj\vivado_ip_prj \vivado_prj.runs\impl_1` folder.

Handshaking Between Preprocessing DUT and Deep Learning IP Core

The DL IP core expects the preprocessed data to be at a specific address in the DDR memory and to have a specific size. The handshaking between the Preprocessing DUT and the DL IP core is to convey the expected address and size to the Preprocessing DUT. The handshaking comprises these steps:

1 The Preprocessing DUT drives the `rd_addr`, `rd_len`, and `rd_avalid` control signals in the `AXIReadCtrlOutDL` bus.

2 The DL IP core samples these control signals and responds to the Preprocessing DUT by sending the data at the `rd_addr` location through the `AXIReadDataDL` signal. The DL IP core also drives the corresponding control signals, `rd_dvalid` and `rd_aready`, in the `AXIReadCtrlInDL` bus.

3 This process continues for three different addresses corresponding to `InputValid` (x"354"), `InputAddr` (x"358"), and `InputSize` (x"35C") signals. The IP core generation report for the DL IP contains the addresses for these registers.
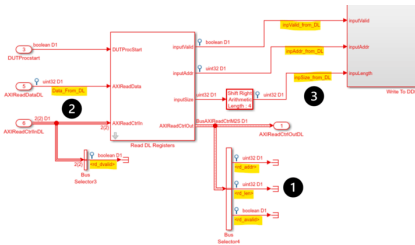
Signals Required for Debugging

The `DLHandshakeLogicExtMem` model contains these signals.

- **rd_addr** --- Address location in the DL IP from which the Preprocessing DUT fetches the required information during handshaking.

- **rd_len** --- Size of data, in bytes, to read from the DL IP starting from the `rd_addr` address location.

- **rd_avalid** --- Indication of whether the data in the `rd_addr` and `rd_len` signals of the same bus is valid.

- **Data_From_DL** --- Information based on the control information the DL IP receives from the Preprocessing DUT in the `AXIReadCtrlOutDL` bus. The DL IP sends appropriate information on this signal.

- **rd_dvalid** --- Control signal that forms part of the `AXIReadCtrlInDL` bus. This signal validates the data in the `AXIReadDataDL` signal.

- **inputAddr_from_DL** --- Output of the `Read DL Registers` subsystem. The Preprocessing DUT places the preprocessed data in the DDR memory at this address.

- **inputSize_from_DL** --- Output of the `Read DL Registers` subsystem. This output is the size of the data that the Preprocessing DUT places in the DDR memory.

- **inputValid_from_DL** --- Output of the `Read DL Registers` subsystem. This signal validates the data in the `inputAddr_from_DL` and `inputSize_from_DL` signals.

Timing Diagram

This timing diagram shows the sequence of events for this scenario.
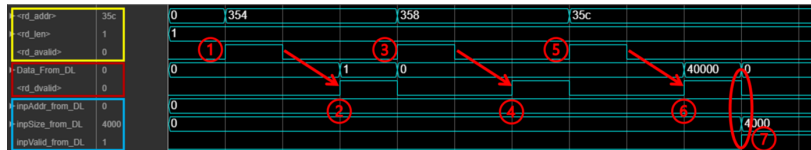
Trigger Conditions in FPGA Data Capture

A successful handshaking between the Preprocessing DUT and DL IP comprises seven events. These events act as sequential triggers in the **FPGA Data Capture** tool to capture the data.
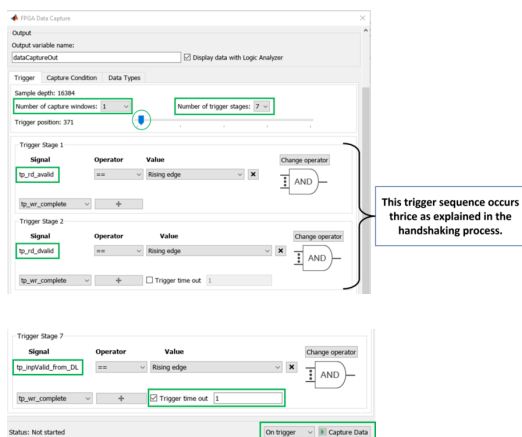
Configure these settings in the **FPGA Data Capture** tool:

- Set **Number of capture windows** to 1 to indicate that handshaking events happen only at the beginning of preprocessing. The signal data corresponding to the entire sample depth can be captured in a single window once these trigger conditions are satisfied.
- Set **Number of trigger stages** to 7 to indicate that the handshaking comprises seven events.
- Set **Trigger Position** to a small value close to zero. If you set this option to 0, you cannot visualize these events because the tool captures signal data only after this trigger.
- Repeat the **Trigger Stage 1** and **Trigger Stage 2** sequences three times.
- Use a trigger time out to ensure that **Trigger Stage 7** happens within one clock cycle of **Trigger Stage 6**. **Trigger Stage 7** corresponds to a rising edge on the `inpValid_from_DL` signal
- Set **Capture mode** to `On Trigger`.

Visualize Captured Data in Logic Analyzer

This timing diagram shows that the handshaking between the Preprocessing DUT and the DL IP behaves as expected.
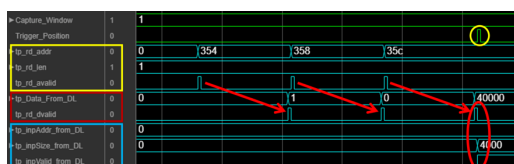
Functionality of `Resize` Subsystem

In this scenario, the focus is to verify the behavior of the `Resize` subsystem. The input image to the `Resize` subsystem is of size 224-by-340 (76,160 pixels). The output image of the `Resize` subsystem is of size 128-by-128 (16,384 pixels). You can use FPGA data capture feature to count the total number of output pixels from the `Resize` subsystem and capture the resized image data to find any errors within the logic. Simulink™ does not support renaming of the output of a Bus Selector block. To rename the signal, use the model components contained in the green boxes in this image.
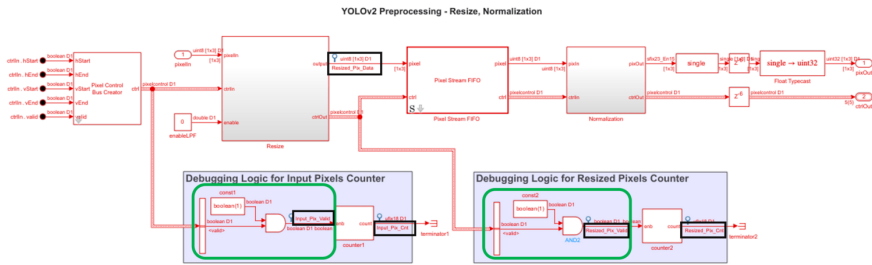
Signals Required for Debugging

The `YOLOv2PreprocessAlgoDebug` model contains these signals.

- **Input_Pix_Valid** --- Control signal that is a part of the `pixelcontrol` bus input of the `Resize` subsystem. This signal validates the pixel data in the `Inp_Pixel_Data` signal.

- **Input_Pix_Cnt** --- Output of the HDL Counter block, which counts the number of valid pixels that you pass as input to the `Resize` subsystem. The model uses the `Input_Pix_Valid` signal to enable this counter.

- **Resized_Pix_Data** --- Output signal of the `Resize` subsystem. This signal contains the pixel data corresponding to the resized image.

- **Resized_Pix_Valid** --- Control signal that is a part of the `pixelcontrol` bus output of the `Resize` subsystem. This signal validates the pixel data in the `Resized_Pix_Data` signal.

- **Resized_Pix_Cnt** --- Output of the HDL Counter block, which counts the number of valid pixels returned by the `Resize` subsystem. The model uses the `Resized_Pix_Valid` signal to enable this counter.

Timing Diagram

Validate the output pixel data using the `Resized_Pix_Valid` signal. Whenever this signal goes high, the `Resize` subsystem sends the valid output data, as this timing diagram shows. The `Input_Pix_Cnt` and `Resized_Pix_Cnt` signals indicate the number of valid pixels entering and emerging from the `Resize` subsystem, respectively.
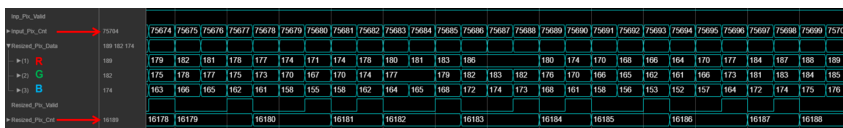
Trigger Conditions in FPGA Data Capture

To capture the valid resized pixel data, use the capture condition logic in the FPGA Data Capture tool.
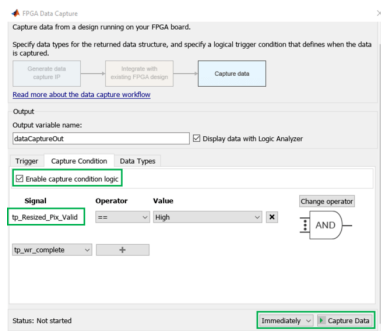
Configure these settings in the **FPGA Data Capture** tool:

- Select **Enable the capture control logic** in the **Capture Condition** tab.
- Use the `Resized_Pix_Valid` signal in the capture condition logic to ensure that the tool captures the data only when this signal goes high.
- Select `Immediately` in the capture mode dropdown menu to enable immediate capture. This option is suitable for scenarios in which no specific triggers determine when the tool captures data.

Visualize Captured Data in Logic Analyzer

This timing diagram shows the resized pixel data and the pixel counts captured by the FPGA Data Capture tool. The `tp_Resized_Pix_Valid` signal is always high, unlike in the equivalent model simulations using Simulink software. This discrepancy is because the capture condition indicates that the FPGA Data Capture tool captures data only when `tp_Resized_Pix_Valid` is high.

The **FPGA Data Capture** tool creates the `dataCaptureOut` structure in the MATLAB workspace after it captures data. Visualize the resized image by extracting and concatenating the RGB image data from `dataCaptureOut`.

```
RData = reshape(dataCaptureOut.tp_Resized_Pix_Data_0,128,128);
BData = reshape(dataCaptureOut.tp_Resized_Pix_Data_2,128,128);
GData = reshape(dataCaptureOut.tp_Resized_Pix_Data_1,128,128);
resizedImage = cat(3,RData',GData',BData');
imshow(resizedImage)
```
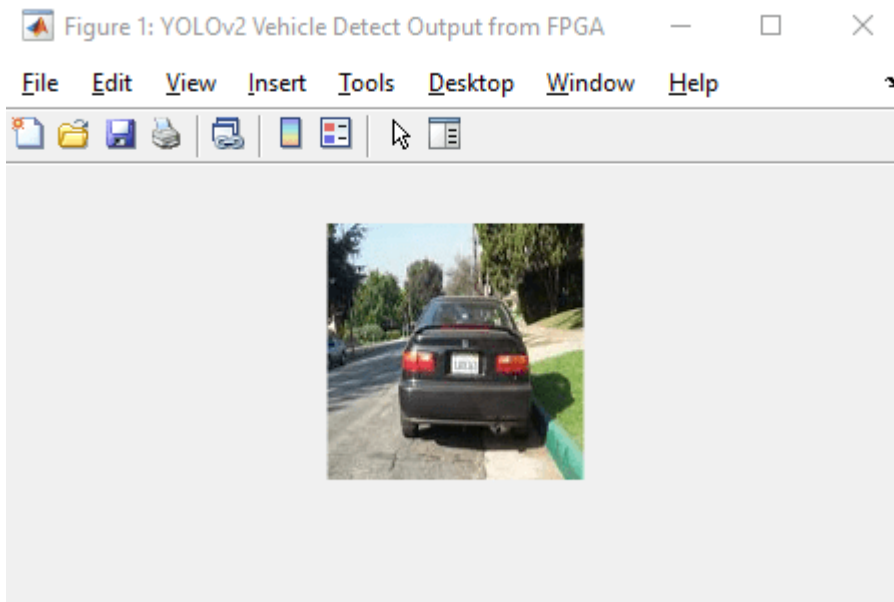
Scenario 3: Handshaking Between Preprocessing DUT and DDR Memory

After the Preprocessing DUT resizes and normalizes the input image, it places the preprocessed image data in the DDR memory at the address it receives from the DL IP. The handshaking process comprises these steps:

1   The Preprocessing DUT drives the `wr_addr`, `wr_len`, and `wr_valid` control signals in the `AXIWriteCtrlOutDDR` bus. The DUT also sends the preprocessed signal data through the `AXIWriteDataDDR` signal.

2   The DDR memory samples these control signals and the preprocessed pixel data received from the Preprocessing DUT.

3   Once all the data is placed in the DDR memory, the DDR memory acknowledges the Preprocessing DUT with a pulse on the `wr_complete` signal in the `AXIWriteCtrlInDDR` bus.
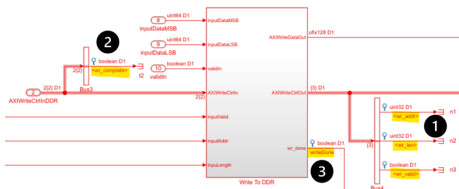
Signals Required for Debugging

The `DLHandshakeLogicDebug` model contains these signals.

*   **wr_addr** --- Control signal that is a part of the `AXIWriteDataDDR` bus. This signal is the address in the DDR memory at which the Preprocessing DUT places the data.

*   **wr_len** --- Control signal that is a part of the `AXIWriteDataDDR` bus. This signal is the size of data, in bytes, that the Preprocessing DUT places in the DDR memory starting from the `wr_addr` address location.

*   **wr_valid** --- Control signal that is a part of the `AXIWriteDataDDR` bus. This signal validates the data in the `wr_addr`, and `wr_len` signals of the same bus.

- **wr_complete** --- Control signal that is a part of the `AXIWriteCtrlInDDR` bus. This signal is the acknowledgement sent from the DDR memory to the Preprocessing DUT containing an indication of the status of the data.
- **writeDone** --- Output of the `Write To DDR` subsystem. This signal indicates whether the data transfer to the DDR memory is successful and triggers the DL IP to start reading that data from the DDR memory for further processing.

Timing Diagram

After the final rising edge on the `wr_valid` control signal occurs, the DDR memory sends a pulse on the `wr_complete` signal as an acknowledgement and a pulse sent on the `writeDone` internal signal. This timing diagram shows the sequence of events for this scenario.
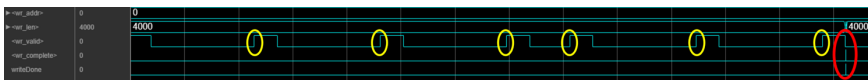
Trigger Conditions in FPGA Data Capture

Configure these settings in the **FPGA Data Capture** tool:

- Set **Number of capture windows** to `1` because these handshaking events happen towards the end of the transaction between Preprocessing DUT and the DDR memory. After these trigger conditions are satisfied, the signal data corresponding to the entire sample depth can be captured in a single window.
- Set **Number of trigger stages** to `2` because this handshaking event comprises three events, of which two events occur simultaneously.
- Set **Trigger position** option close to the end of the handshake to ensure the Logic Analyzer displays the complete handshake.
- Set **Capture mode** to `On Trigger`.

The **Trigger Stage 1** corresponds to a rising edge on `wr_valid` signal from the DDR memory.

The **Trigger Condition 2** section captures an expected pulse on the `wr_complete` and `writeDone` signals. This stage uses logical and comparison operators.
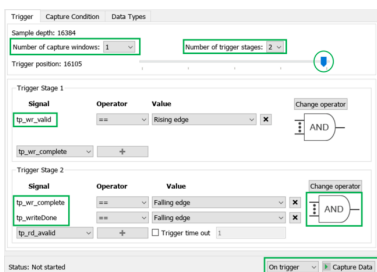
Visualize Captured Data in Logic Analyzer

This timing diagram confirms that the handshaking between Preprocessing DUT and DDR memory happens as expected.
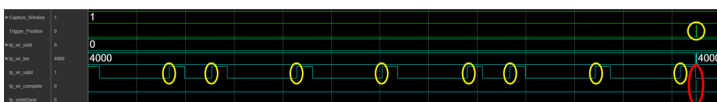
Use FPGA Data Capture and AXI Manager Features Simultaneously

As described in "Design Considerations for Data Capture" (HDL Verifier), to use AXI manager and FPGA data capture features simultaneously, set the capture mode of FPGA data capture to nonblocking. Create an FPGADataCapture object in non-blocking mode and launch the **FPGA Data Capture** tool.

```
cd(fullfile('hdl_prj','ipcore','YOLOV2Pre_cs_ipv4_v1_0','fpga_data_capture'))
fpgadc = FPGADataCapture;
fpgadc.CaptureMode = 'nonblocking';
launchApp(fpgadc);
```

You must configure a few registers before sending a video frame as an input to the model. Set the `DUTProcStart` register of the Preprocessing DUT to `1`. AXI manager can be leveraged to do this task. The `YOLOv2DeployAndVerifyDetector` function that is attached with the "Deploy and Verify YOLO v2 Vehicle Detector on FPGA" on page 3-14 example has all the steps present in Verify Deployed YOLO v2 Vehicle Detector Using MATLAB section. The `YOLOv2DeployAndVerifyDetector` function uses `writePort` function to configure all the control registers. To use the AXI manager instead of `writePort` to configure the `DUTProcStart` register, use the `helperUpdateYOLOv2DeployAndVerifyDetector` function.

The `helperUpdateYOLOv2DeployAndVerifyDetector` function creates the `DebugYOLOv2VehicleDetector` function which is a modified version of the `YOLOv2DeployAndVerifyDetector` function and contains an object of the AXI manager. The `helperUpdateYOLOv2DeployAndVerifyDetector` function adds this code to the `DebugYOLOv2VehicleDetector` function, which you can use to access AXI manager feature.

Create an AXI manager object.

```
h = aximanager('Xilinx');
```

Use `writememory` function to write `1` into the `DUTProcStart` register. The address for this register can be found in the IP Core Generation report.

```
writememory(h, '0xA0040100',1);
```

Release the JTAG cable resource after writing into the `DUTProcStart` register to ensure that FPGA data capture can use the same JTAG interface to capture the data.

```
release(h)
```

To capture the required data corresponding to different scenarios, the **FPGA Data Capture** tool with the appropriate trigger conditions. This diagram shows the data capture process:
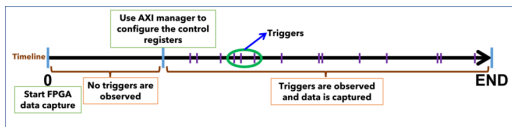
1 Configure the **FPGA Data Capture** tool with the trigger conditions and then click the **Capture Data** button to start the data capture process. The tool captures the data when it observes triggers.

2 Enter the command `DebugYOLOv2VehicleDetector(hSOC)` to start the workflow comprising all the steps from configuring the registers to reading back the processed data to MATLAB.

Because you start the FPGA Data Capture tool before this step, the FPGA Data Capture tool detects all the events.

The AXI manager configures the `DUTProcStart` control register while the FPGA Data Capture tool waits for the trigger condition to be satisfied. You can simultaneously use both of these tools to capture all the required data.

Conclusions

In summary, this example shows how to instrument a Simulink model with debug hooks to allow visibility of signals after deploying your design to an FPGA or SoC board. You use AXI manager to configure the control registers in the deployed design from MATLAB and then specify the triggers in the **FPGA Data Capture** tool for capturing the signals of interest. You analyze the captured data and use the results to debug your application.

See Also

"Deploy and Verify YOLO v2 Vehicle Detector on FPGA" on page 3-14

More About

- "Data Capture Workflow" (HDL Verifier)
- "Set Up AXI Manager" (HDL Verifier)
- "Target Deep Learning Processor and Image Preprocessing to FPGA" (Vision HDL Toolbox Support Package for Xilinx Zynq-Based Hardware)

Integrate YOLO v2 Vehicle Detector System on SoC

This example shows how to simulate a you only look once (YOLO) vehicle detector and verify the functionality of the end-to-end application using MATLAB.

The end-to-end application includes preprocessing of the images, a YOLO v2 vehicle detection network, and postprocessing of the images to overlay results.

Load Camera Data and Network File

This example uses a PandasetCameraData.mp4 file that contains a subset of the video from PandaSet data set. Download the video file and the network .mat file.

```
supportFileDir =  matlab.internal.examples.utils.getSupportFileDir();
pathToDataset = fullfile(supportFileDir, 'visionhdl', 'PandasetCameraData');
if(~isfile(fullfile(pathToDataset, 'PandasetCameraData.mp4')) ...
                || ~isfile(fullfile(pathToDataset, 'yolov2VehicleDetector32Layer.mat')) ...
                || ~isfile(fullfile(pathToDataset, 'yolov2VehicleDetector60Layer.mat')))
    PandasetZipFile = matlab.internal.examples.downloadSupportFile('visionhdl','PandasetCameraDat
    [outputFolder,~,~] = fileparts(PandasetZipFile);
    unzip(PandasetZipFile,outputFolder);
end
```

```
addpath(pathToDataset);
```

A YOLO v2 vehicle detection application has three main modules. The preprocessing module, accepts the input frame and performs image resize and normalization. The preprocessed data is then consumed by the YOLO v2 vehicle detection network, which is a feature extraction network followed by a detection network. The network output is postprocessed to identify the strongest bounding boxes and the resulting bounding boxes are overlaid on the input image.
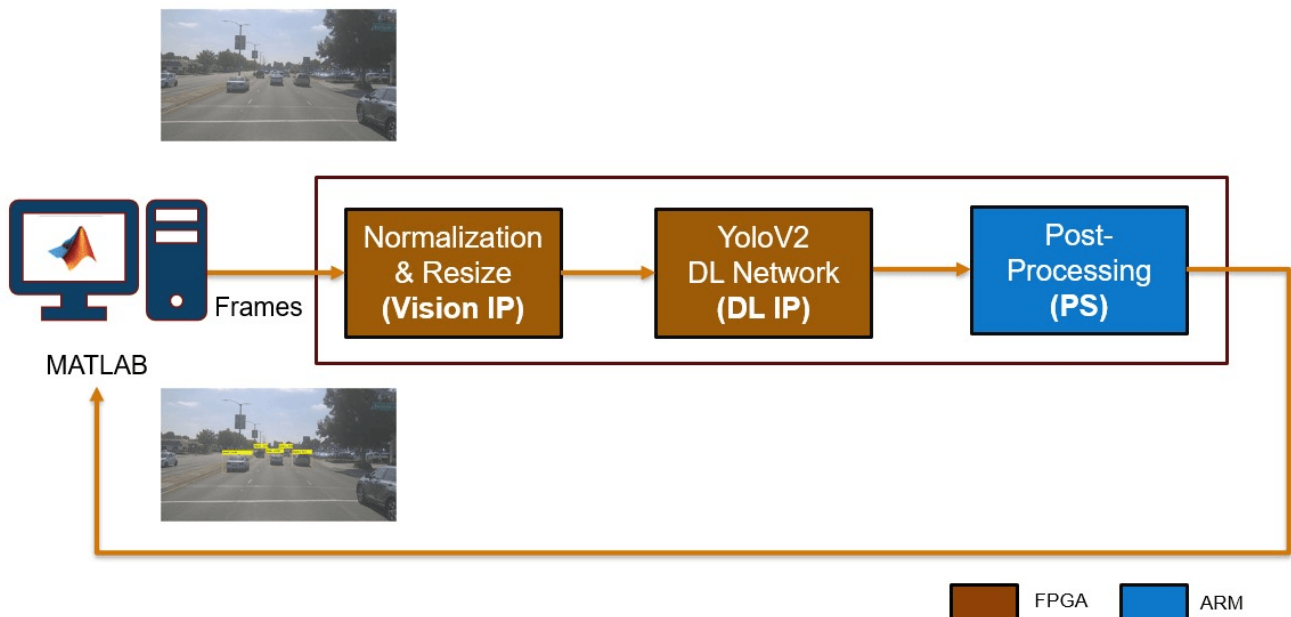
The preprocessing subsystem and DLIP are deployed on FPGA (Programmable Logic, PL) and the postprocessing is deployed on the ARM prcoessor (Processing System, PS). For deploying the vehicle detector, see "YOLO v2 Vehicle Detector with Live Camera Input on Zynq-Based Hardware" on page 3-49. This example shows how to model the preprocessing module (resize and normalization) and postprocessing module along with DL handshaking logic and network execution.

Explore Vehicle Detector

```
open_system('YOLOv2VehicleDetectorOnSoC');
```


YOLO v2 vehicle detector on SoC - Preprocess and Postprocess with deep learning handshake logic
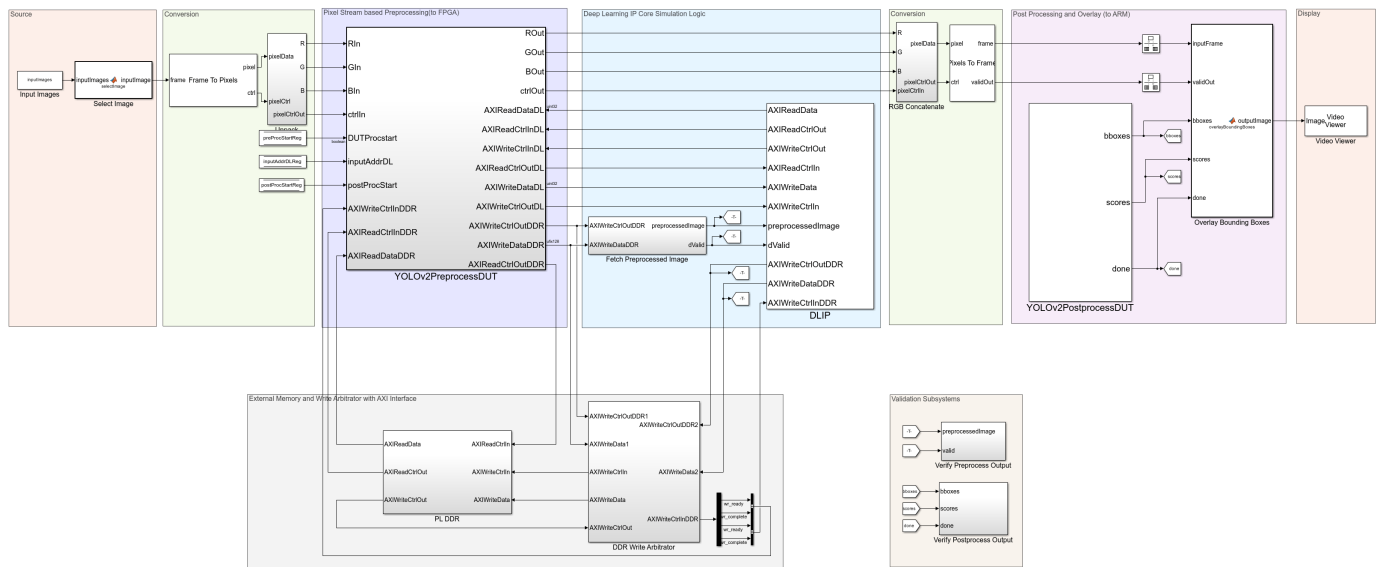
Copyright 2022-2023 The MathWorks, Inc.

The vehicle detector contains these modules:

- Source- Selects the inputImage from Pandaset.
- Conversion- Converts the input frame into RGB pixel stream.
- Pixel-stream based preprocessing(to FPGA)- Preprocesses the input frame and writes it into DDR.
- Deep learning IP Core Simulation Logic- Models the DL processor to calculate activations on the input frame and write the output to DDR
- Conversion- Converts the input RGB pixel stream to frame for overlaying bounding boxes.
- Postprocessing and Overlay(to ARM)- Applies postprocessing to network output and overlay the bounding boxes on the input frame.
- Display- Displays the input frame with detections.

The `inputImages` stores `numFrames` number of images from Pandaset. The frame is initially resized and normalized in `YOLOv2PreprocessDUT` and the preprocessed output is written into DDR at the address location read from DL input handshaking registers, (`InputValid`, `InputAddr`, `InputSize`). The `DLIP` calculates activations on the preprocessed image, writes the activations to DDR, and updates the DL output handshaking registers, (`OutputValid`, `OutputAddr`, `OutputSize`). This handshaking triggers the `YOLOv2PostprocessDUT`, that reads the DL output

from the address information obtained from the DL registers, and performs post processing and calculates bounding boxes that are displayed in the `VideoViewer` block via the `overlayBoundingboxes` function.

YOLOv2PreprocessDUT

open_system('YOLOv2VehicleDetectorOnSoC/YOLOv2PreprocessDUT');
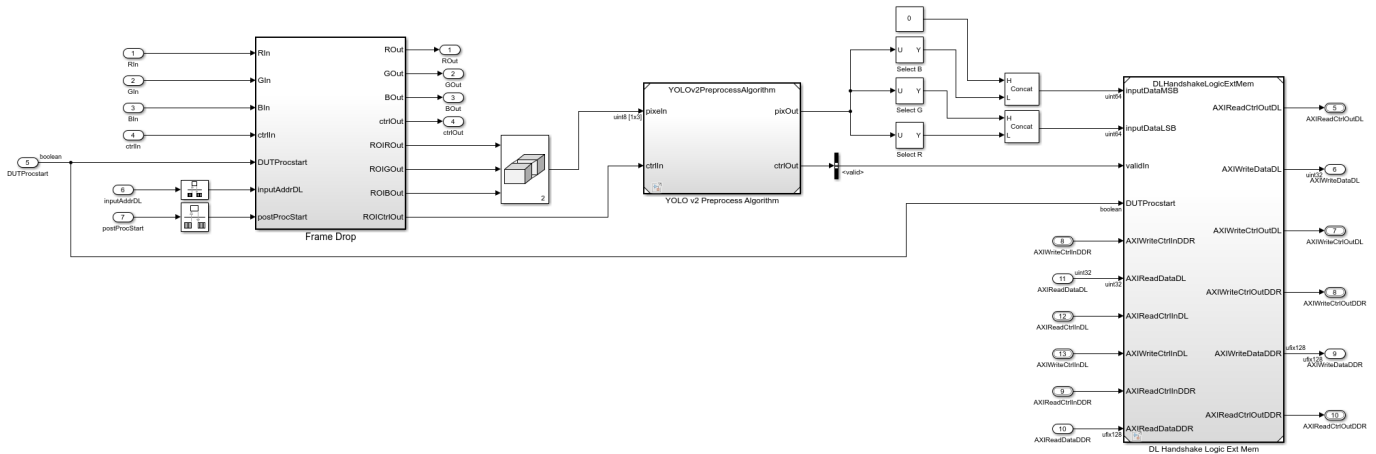
The `selectImage` subsystem selects the input frame from `inputImages` block. A `Frame To Pixels` block converts the input image from the `selectImage` to a pixel stream and `pixelcontrol` bus. The `Unpack` subsystem divides the pixel stream into R, G, B components. The RGB data, (`RIn`, `GIn`, `BIn`) along with ctrl bus is fed for preprocessing. The input image is streamed out as, (`ROut`, `GOut`, `BOut`) to write it into the PS DDR for overlaying the bounding boxes.

The `YOLOv2PreprocessDUT` contains subsystems for frame dropping, selecting Region of Interest (ROI) from the input frame, preprocessing (resize and normalization), and handshaking logic.

The `Frame Drop` subsystem synchronizes data between `YOLOv2PreprocessDUT` and `DLIP` by dropping the input frames if `DLIP` is not available for processing. It contains finite state machine (FSM) logic for reading `DLIP` registers and a pixel bus creator to concatenate the output control signals of frame drop logic to `pixel control` bus. The `readInputRegisters` subsystem reads the `inputAddrReg` register and forwards the first frame to preprocessing and resets the control signals for rest of the frames until `inputAddr` is updated by DLIP. This frame drop logic lets the DLIP process one frame corresponding to one `inputAddr`.

The output of the `Frame Drop` subsystem is sent to the `ROI Selector` block that selects the ROI from the input image and forwards it for preprocessing. The ROI is selected for the input image from Pandaset of size 1920x1080 and is scaled down by a factor of 4 for faster simulation. The ROI is configured in `helperSLYOLOv2SimulationSetup` function.

```
hPos = 350;
vPos = 400;
hSize = 1000;
vSize = 600;
```

The `YOLO v2 Preprocess Algorithm` contains subsystems to perform resizing and normalization operations. The pixel stream from the `Frame Drop` subsystem is passed to the `Resize` subsystem for resizing the input image to the input size expected by the deep learning network, (`128, 128, 3`). The resized output is passed to `Normalization` subsystem for rescaling the pixel values to [0, 1]

range. This preprocessed frame is then passed to the `DL Handshake Logic Ext Mem` subsystem to be written into the PL DDR.

The `DL Handshake Logic Ext Mem` subsystem contains a finite state machine (FSM) logic for handshaking with `DLIP` and a subsystem to write the frame to DDR. The `Read DL Registers` subsystem has the FSM logic to read the handshaking signals (`InputValid, InputAddr, and InputSize`) from the `DLIP` for multiple frames. The `Write to DDR` subsystem uses these handshaking signals to write the preprocessed frame to the memory using AXI4-Master protocol. For more information on the `Yolov2PreprocessDUT` refer to the example, "Deploy and Verify YOLO v2 Vehicle Detector on FPGA" on page 3-14

DLIP

```
open_system('YOLOv2VehicleDetectorOnSoC/DLIP','force');
```
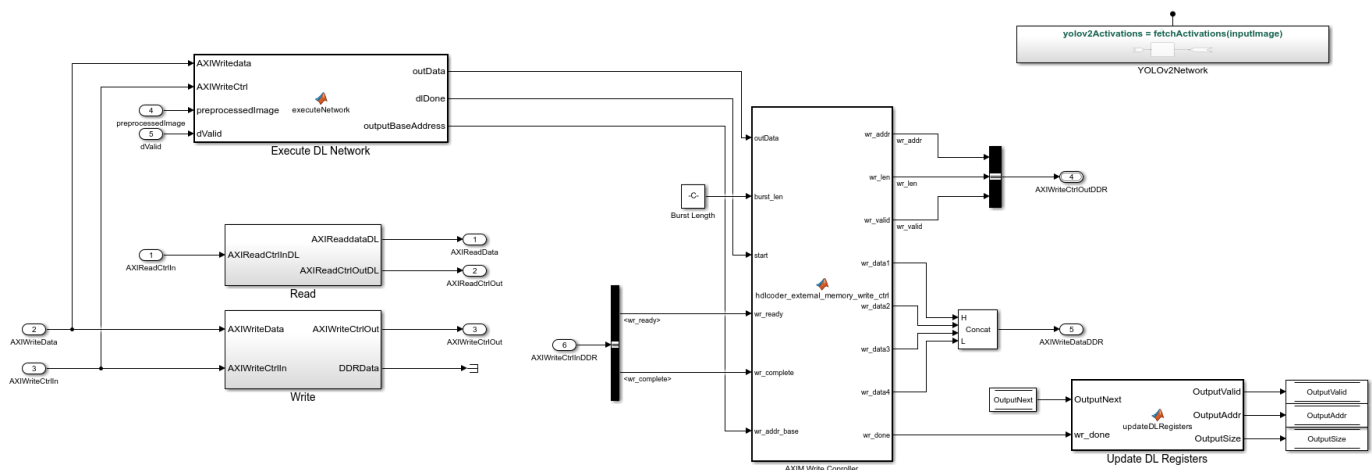

The `DLIP` contains subsystems for prediction logic, DL input and output register handshaking logic, and an AXI Write controller to write the DL Output to DDR.

The `FetchPreprocessedImage` subsystem reads and rearranges the output from `YOLOv2PreprocessDUT` to the `networkInputSize` as required by the deep learning network. The network and the activation layer of the `DLIP` are setup using `helperSLYOLOv2SimulationSetup` and `helperYOLOv2Network` functions.

This example uses a pretrained YOLO v2 network that was trained on Pandaset. The network output is rearranged to the external memory data format of the DL Processor by concatenating the elements along the third dimension. For more information, see "External Memory Data Format" (Deep Learning HDL Toolbox).

The DL output is written to memory using `AXIM Write Controller` subsystem. The write operations from the `YOLOv2PreprocessDUT` and `DLIP` are multiplexed using `DDR Write Arbitrator`.

YOLOv2PostprocessDUT

```
open_system('YOLOv2VehicleDetectorOnSoC/YOLOv2PostprocessDUT','force');
```
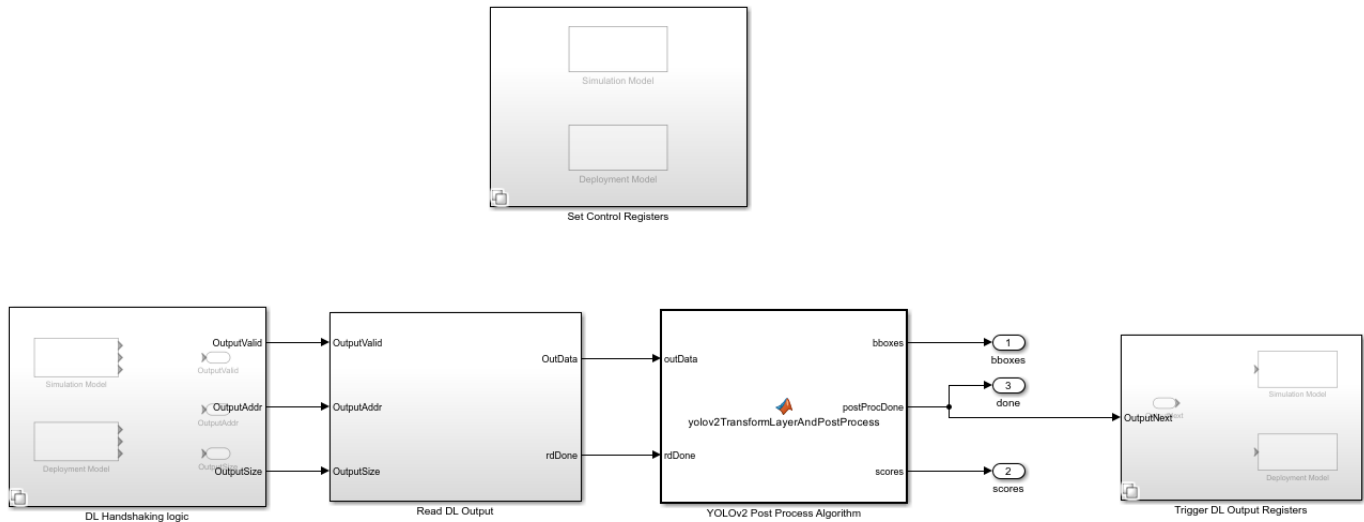
The `YOLOv2PostprocessDUT` subsystem contains subsystems for DL Handshaking, reading DL output, transforming and applying post processing to the DL Output. The DL handshaking subsystems have variant behavior depending on whether the model is configured for simulation or deployment based on `simulationFlag`. Since this example demonstrates the simulation workflow, the `simulationFlag` is set to true in `helperSLYOLOv2Setup` script.

The `Set Control Registers` subsystem sets the control registers for `YOLOv2PreprocessDUT`, `postProcStart`, `DUTProcStart`, and `inputAddrReg`. The `DL Handshaking` subsystem reads the DL Output handshaking registers, (`OutputValid`, `OutputAddr`, `OutputSize`) indicating address, size, and validity of the output. The model abstracts these registers as datastore blocks for simulation. The `readDLOutput` subsystem uses these handshaking signals and reads the DL Output from PL DDR.

The `readDLOutput` subsystem contains subsystems for polling `OutputValid`, generating read requests, and reading DL output from PL DDR. The `pollOutputValid` function polls for the `OutputValid` signal from `DLIP` and triggers post processing when `OutputValid` is asserted. The `read DL Output from PL DDR` subsystem contains a signal `rdDone` which indicates that DL Output read operation is completed successfully. The `TriggerDLOutputNext` subsystem pulses `OutputNext` signal when `rdDone` is asserted to indicate to the `DLIP` that the output of current frame is read.

The DL output data is then sent to `yolov2TransformlayerandPostprocess` function for postprocessing. It transforms the DL Output from DDR by rearranging, normalizing the data, and thresholding the bounding boxes with a confidence score of 0.4. It returns the bounding boxes and pulses `postProcDone` signal to indicate that the post processing is completed successfully.

The `YOLOv2PostprocessDUT` is configured with these DL network parameters, `networkInputSize`, `networkOutputSize`, `anchorBoxes` and `inputImageROI`, `inputROISize`, `confidenceThreshold` in `helperSLYOLOv2SimulationSetup.m` script.

```
vehicleDetector = load(networkmatfile);
detector = vehicleDetector.detector;
net = detector.Network;
anchorBoxes = detector.AnchorBoxes;
networkInputSize = net.Layers(1, 1).InputSize;
networkOutputSize = [16,16,12];
```

```
paddedOutputSize = (networkOutputSize(1)*networkOutputSize(2)*networkOutputSize(3)*4)/3;
inputImageROI = [hPos, vPos, hSize, vSize];
inputROISize = [vSize, hSize, numComponents];
confidenceThreshold = 0.4;
```

Simulate Vehicle Detector

Configure the network for the vehicle detector using the `helperSLYOLOv2SimulationSetup` function.

```
helperSLYOLOv2SimulationSetup();
```

The script supports 2 networks, a 32 layer network(default) and a 60 layer network. To run the 60 layer network, set the networkConfig to '60layer'.

```
helperSLYOLOv2SimulationSetup('60layer');
```

This model takes a couple of minutes to update the diagram when you are compiling for the first time. Update the model before running the simulation.

```
set_param("YOLOv2VehicleDetectorOnSoC", SimulationCommand="update");
out = sim("YOLOv2VehicleDetectorOnSoC");
```

```
### Starting serial model reference simulation build.
### Model reference simulation target for DLHandshakeLogicExtMem is up to date.
### Model reference simulation target for YOLOv2PreprocessAlgorithm is up to date.

Build Summary

0 of 2 models built (2 models already up to date)
Build duration: 0h 0m 32.939s
```
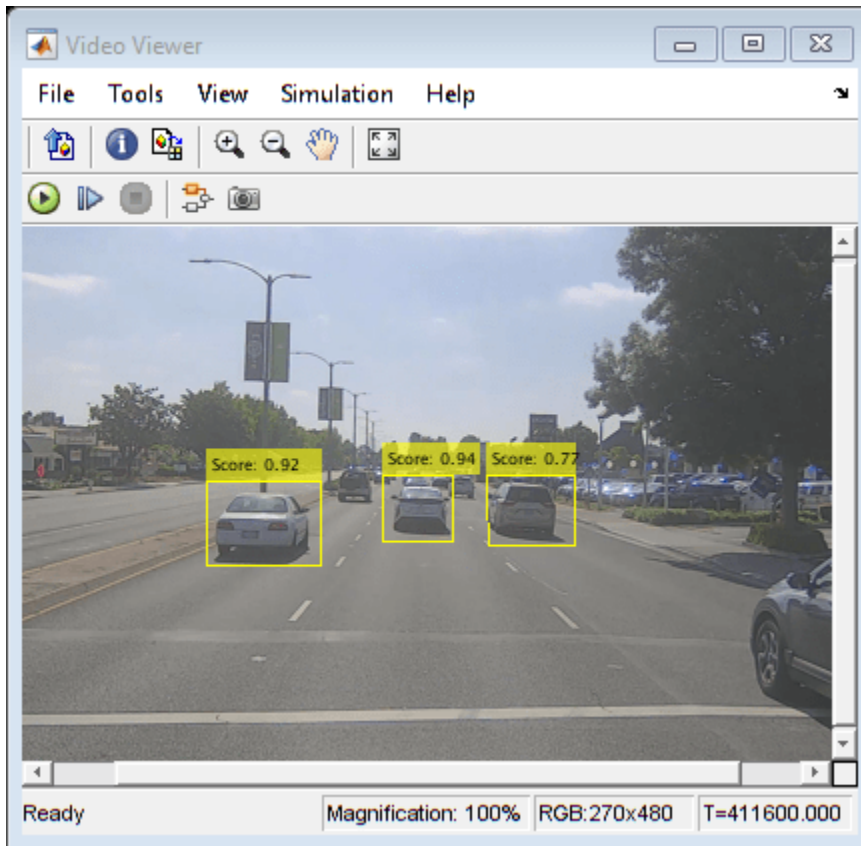
Verify YOLOv2PreprocessDUT and YOLOv2PostprocessDUT using MATLAB

The example includes subsystems for verification of outputs of `YOLOv2PreprocessDUT` and `YOLOv2PostprocessDUT`. The `Verify Preprocess Output` and `Verify Postprocess Output` subsystems log the signals required for the verification of the preprocessed image and bounding boxes, respectively.

```
helperVerifyVehicleDetector;
```
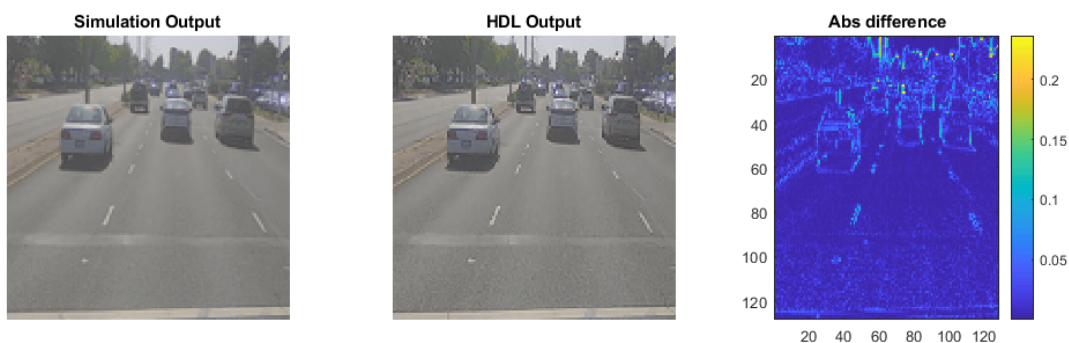
YOLOv2PostprocessDUT Verification
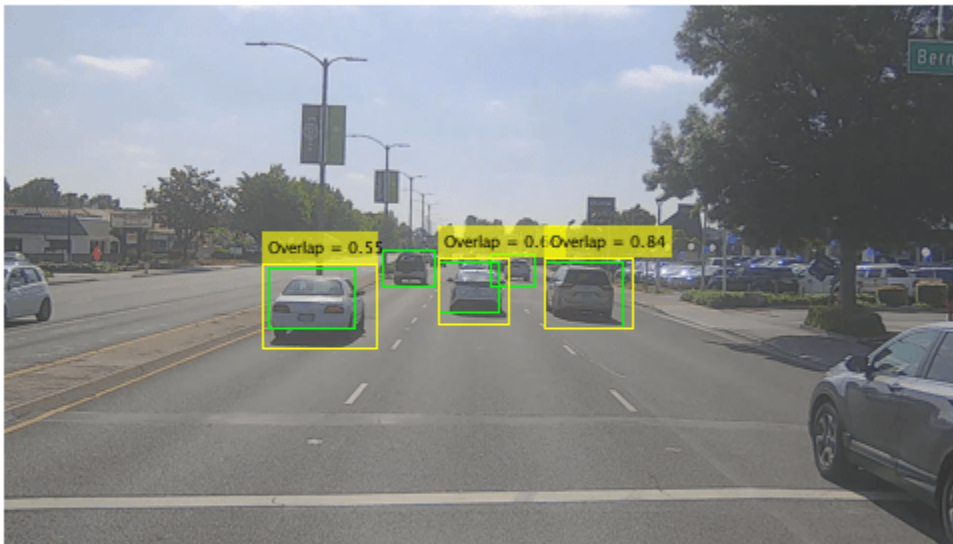

```
--- Simulation
--- YOLOv2PostprocessDUT
```

Close the figures

```
close(hFigurePreprocess);
close(hFigurePostprocess);
```

The `helperVerifyVehicleDetector` script verifies all the logged outputs obtained in simulation. It compares the preprocessed image obtained in simulation with the reference image obtained by applying resize and normalize operations and overlays the bounding boxes obtained from simulation and from `detect` (Computer Vision Toolbox) function on the input images from the dataset.

Conclusion

This example demonstrated the YOLOv2 vehicle detector application comprising of preprocessing steps(image resize and normalization) and handshaking logic on FPGA, vehicle detection using DLIP followed by postprocessing and verified the results using MATLAB.

Copyright 2022-2023 The MathWorks, Inc.

See Also

Related Examples
- "Deploy and Verify YOLO v2 Vehicle Detector on FPGA" on page 3-14

More About
- "Deep Learning Processing of Live Video" (Vision HDL Toolbox Support Package for Xilinx Zynq-Based Hardware)

YOLO v2 Vehicle Detector with Live Camera Input on Zynq-Based Hardware

Introduction

The YOLO v2 Vehicle Detector with Live Camera Input example extends the Deploy and Verify YOLO v2 Vehicle Detector on FPGA example by adding live HDMI video input and by targeting the post processing logic to the ARM processor of the Xilinx® Zynq® UltraScale+™ MPSoC ZCU102 Evaluation Kit. The example uses a new RGB for DL Processor reference design. The reference design passes the HDMI input to the preprocessing logic and also writes the input frame to PS DDR. After preprocessing, the design writes the resized and normalized images to PL DDR where it can be accessed by the DL processor. After the DL processor writes the output back to DDR, the postprocessing code reads the output frames to calculate and overlay bounding boxes. These modified output frames are returned on the HDMI output and can also be accessed in Simulink® by using the Video Capture HDMI block.

Setup Prerequisites

This example follows the algorithm development workflow that is detailed in the "Developing Vision Algorithms for Zynq-Based Hardware" (Vision HDL Toolbox Support Package for Xilinx Zynq-Based Hardware) example. If you have not already done so, please work through that example to gain a better understanding of the required workflow.

If you have not yet done so, run through the guided setup wizard portion of the Zynq support package installation. You might have already completed this step when you installed this support package.

On the MATLAB **Home** tab, in the **Environment** section of the **Toolstrip**, click **Add-Ons > Manage Add-Ons**. Locate *Vision HDL Toolbox Support Package for Xilinx Zynq-Based Hardware*, and click **Setup**.

The guided setup wizard performs a number of initial setup steps, and confirms that the target can boot and that the host and target can communicate.

For more information, see "Setup for Vision Hardware" (Vision HDL Toolbox Support Package for Xilinx Zynq-Based Hardware).

Input Video File and Network

This example uses PandasetCameraData.mp4 created from PandaSet data set as the input and it uses network from yolov2VehicleDetector32Layer.mat file. These files are approximately 47MB and 2MB in size. Download the file from Mathworks support website and unzip the downloaded file.

```
PandasetZipFile = matlab.internal.examples.downloadSupportFile('visionhdl','PandasetCameraDat
[outputFolder,~,~] = fileparts(PandasetZipFile);
unzip(PandasetZipFile,outputFolder);
pandasetVideoFile = fullfile(outputFolder,'PandasetCameraData');
addpath(pandasetVideoFile);
```

Pixel Stream Model Design Under Test

The DUT in this example selects a region of interest (ROI) from the input frames to meet the requirements of the DL processor. The model selects 1000-by-500 region of the incoming 1920-by-1080 video.

Since the DL IP core cannot keep up with the incoming frame rate from the camera, the model also includes frame drop logic. The model only processes frames when the DL processor IP core is ready to accept the data.

Configure Deep Learning Processor and Generate IP Core

The deep learning processor IP core accesses the preprocessed input from the DDR memory, performs the vehicle detection, and loads the output back into the memory. To generate a deep learning processor IP core that has the required interfaces, create a deep learning processor configuration by using the `dlhdl.ProcessorConfig` (Deep Learning HDL Toolbox) class. In the processor configuration, set the `InputRunTimeControl` and `OutputRunTimeControl` parameters. These parameters indicate the interface type for interfacing between the input and output of the deep learning processor. To learn about these parameters, see "Interface with the Deep Learning Processor IP Core" (Deep Learning HDL Toolbox). In this example, the deep learning processor uses the `register` mode for input and output runtime control.

hPC = dlhdl.ProcessorConfig;

hPC.InputRunTimeControl = "register";

hPC.OutputRunTimeControl = "register";

Set the `TargetPlatform` property of the processor configuration object as `Generic Deep Learning Processor`. This option generates a custom generic deep learning processor IP core.

hPC.TargetPlatform = 'Generic Deep Learning Processor';

Use the `setModuleProperty` method to set the properties of the `conv` module of the deep learning processor. These properties can be tuned based on the design choice to ensure that the design fits on the FPGA. To learn more about these parameters, see `setModuleProperty` (Deep Learning HDL Toolbox). In this example, `LRNBlockGeneration` is turned on and `SegmentationBlockGeneration` is turned off to support YOLOv2 vehicle detection network. `ConvThreadNumber` is set to 9.

hPC.setModuleProperty('conv','LRNBlockGeneration', 'on');

hPC.setModuleProperty('conv','SegmentationBlockGeneration', 'off');

hPC.setModuleProperty('conv','ConvThreadNumber',9);

This example uses the Xilinx ZCU102 board to deploy the deep learning processor. Use the `hdlsetuptoolpath` function to add the Xilinx Vivado synthesis tool path to the system path.

hdlsetuptoolpath('ToolName','Xilinx Vivado','ToolPath','C:\Xilinx\Vivado\2022.1\bin\vivado.bat')

Use the `dlhdl.buildProcessor` function with the `hPC` object to generate the deep learning IP core. It takes some time to generate the deep learning processor IP core.

dlhdl.buildProcessor(hPC);

The generated IP core contains a standard set of registers and the generated IP core report. The IP core report is generated in the same folder as ip core with the name `testbench_ip_core_report.html`.

IP core name	dlprocessor
IP core version	1.0
IP core folder	dlhdl_prj\ipcore\dlprocessor_v1_0
IP core zip file name	dlprocessor_v1_0.zip
Target platform	Generic Deep Learning Processor Xilinx
Target tool	Xilinx Vivado
Target language	VHDL
Model	testbench

IP core name and IP core folder are required in a subsequent step in 'Set Target Reference Design' task of the IP core generation workflow of the DUT. The IP core report also has the address map of the registers that are needed for handshaking with input and output of deep learning processor IP core.

Port Name	Port Type	Data Type	Target Platform Interfaces	Interface Mapping
InputNext	Inport	boolean	AXI4	x"350"
OutputNext	Inport	boolean	AXI4	x"360"
StreamingMode	Inport	boolean	AXI4	x"34C"
InputStop	Inport	boolean	AXI4	x"374"
inputStart	Inport	boolean	AXI4	x"224"
FrameCount	Inport	uint32	AXI4	x"24C"
InputValid	Outport	boolean	AXI4	x"354"
InputAddr	Outport	uint32	AXI4	x"358"
InputSize	Outport	uint32	AXI4	x"35C"
OutputValid	Outport	boolean	AXI4	x"364"
OutputAddr	Outport	uint32	AXI4	x"368"
OutputSize	Outport	uint32	AXI4	x"36C"

The registers InputValid, InputAddr, and InputSize contain the values of the corresponding handshaking signals that are required to write the preprocessed frame into DDR memory. The register inputNext is used by the DUT to pulse the inputNext signal after the data is written into memory. These register addresses are setup in the helperSLYOLOv2Setup.m script. The other registers listed in the report are read/written using MATLAB. For more details on interface signals, see the Design Processing Mode Interface Signals section of "Interface with the Deep Learning Processor IP Core" (Deep Learning HDL Toolbox).

Generate and Deploy Bitstream to FPGA OR Target the Algorithm

Use the simulation model from this example "Integrate YOLO v2 Vehicle Detector System on SoC" on page 3-41 for simulation as it uses a reduced input image size and the simulation will be faster. Start the targeting workflow by right clicking the YOLOv2 Preprocessing subsystem in the vzYOLOv2DetectorOnLiveCamera model and selecting HDL Code > HDL Workflow Advisor.

```
open_system('vzYOLOv2DetectorOnLiveCamera');
```

YOLO v2 vehicle detector on SoC - Preprocess and Postprocess with deep learning handshake logic
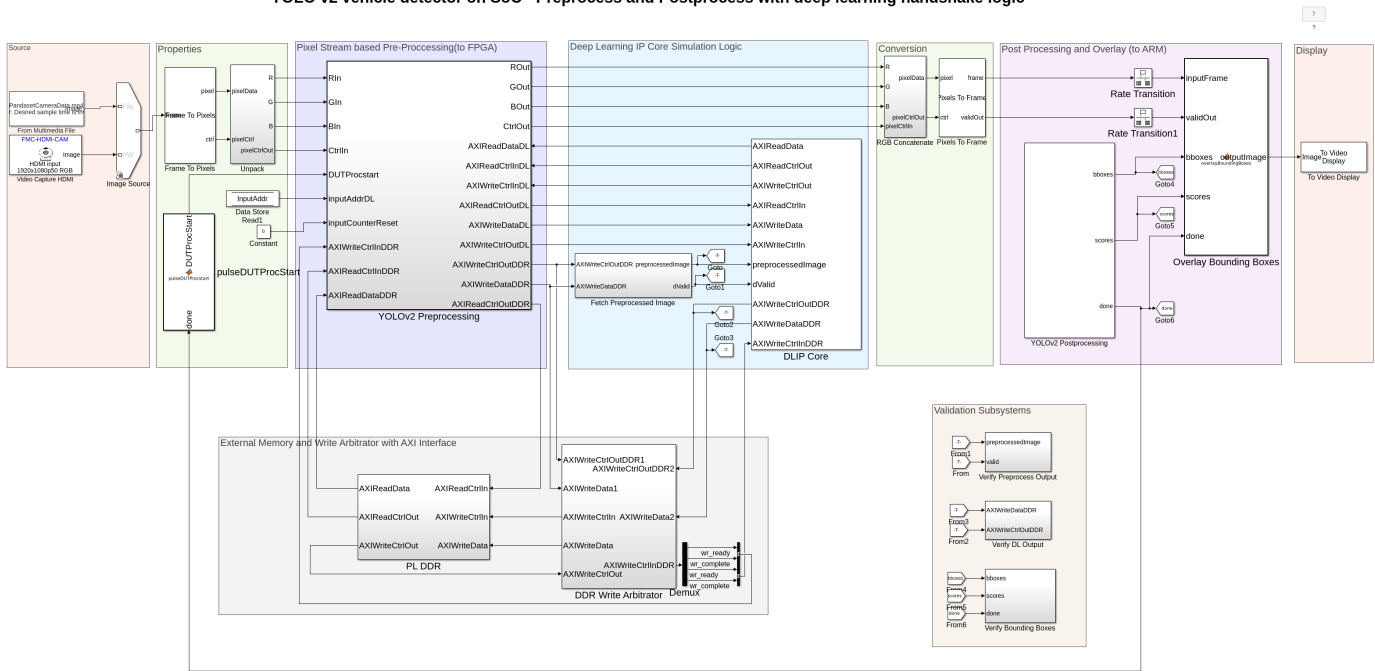
Copyright 2022-23 The MathWorks, Inc.

In step 1.1, select IP Core Generation workflow and the platform 'ZCU102 with FMC-HDMI-CAM'.
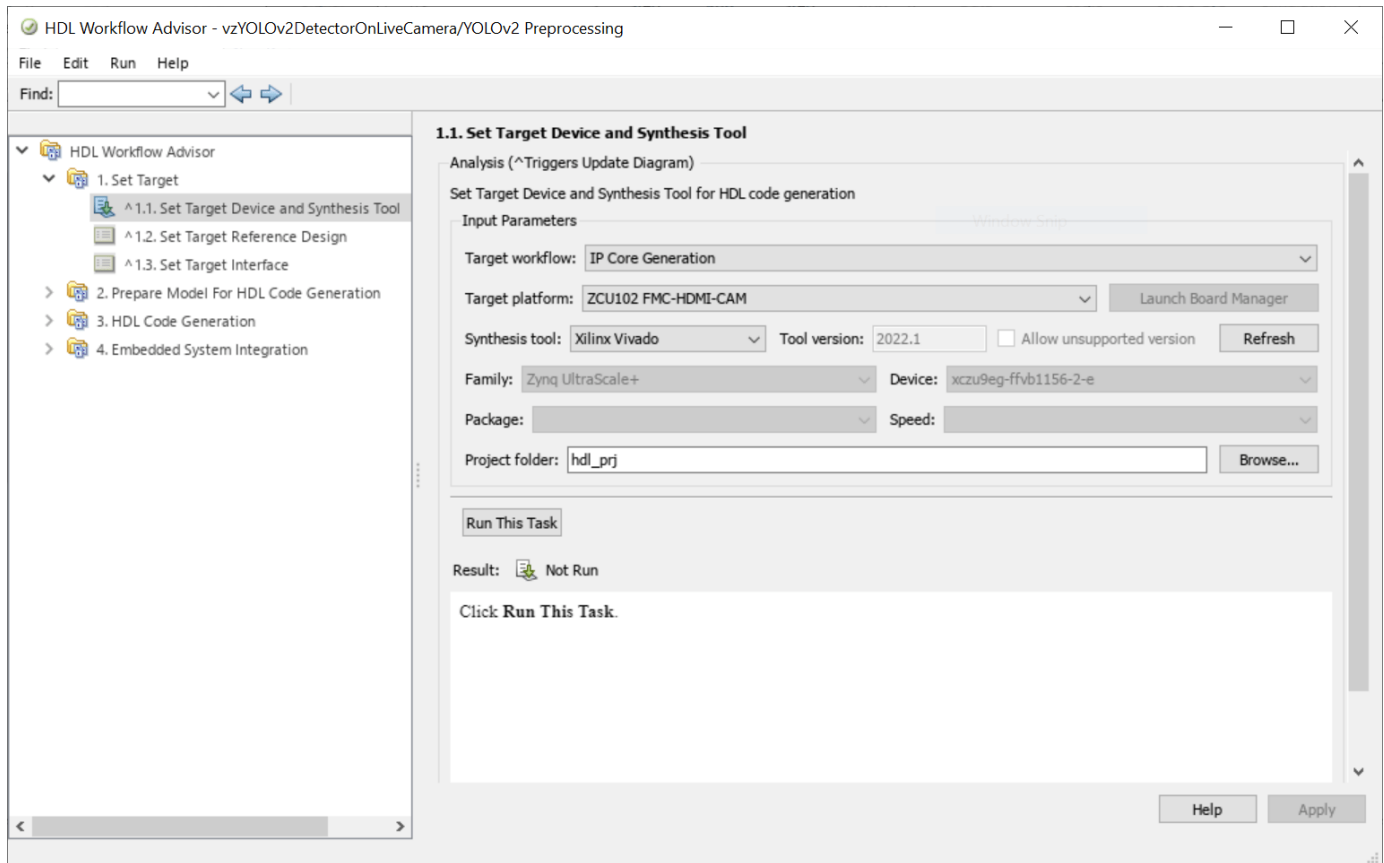
In step 1.2, the reference design is set to "RGB with DL Processor". The DL Processor IP name and the DL Processor IP location specify the name and location of the generated deep learning processor IP core, and are obtained from the IP core report.
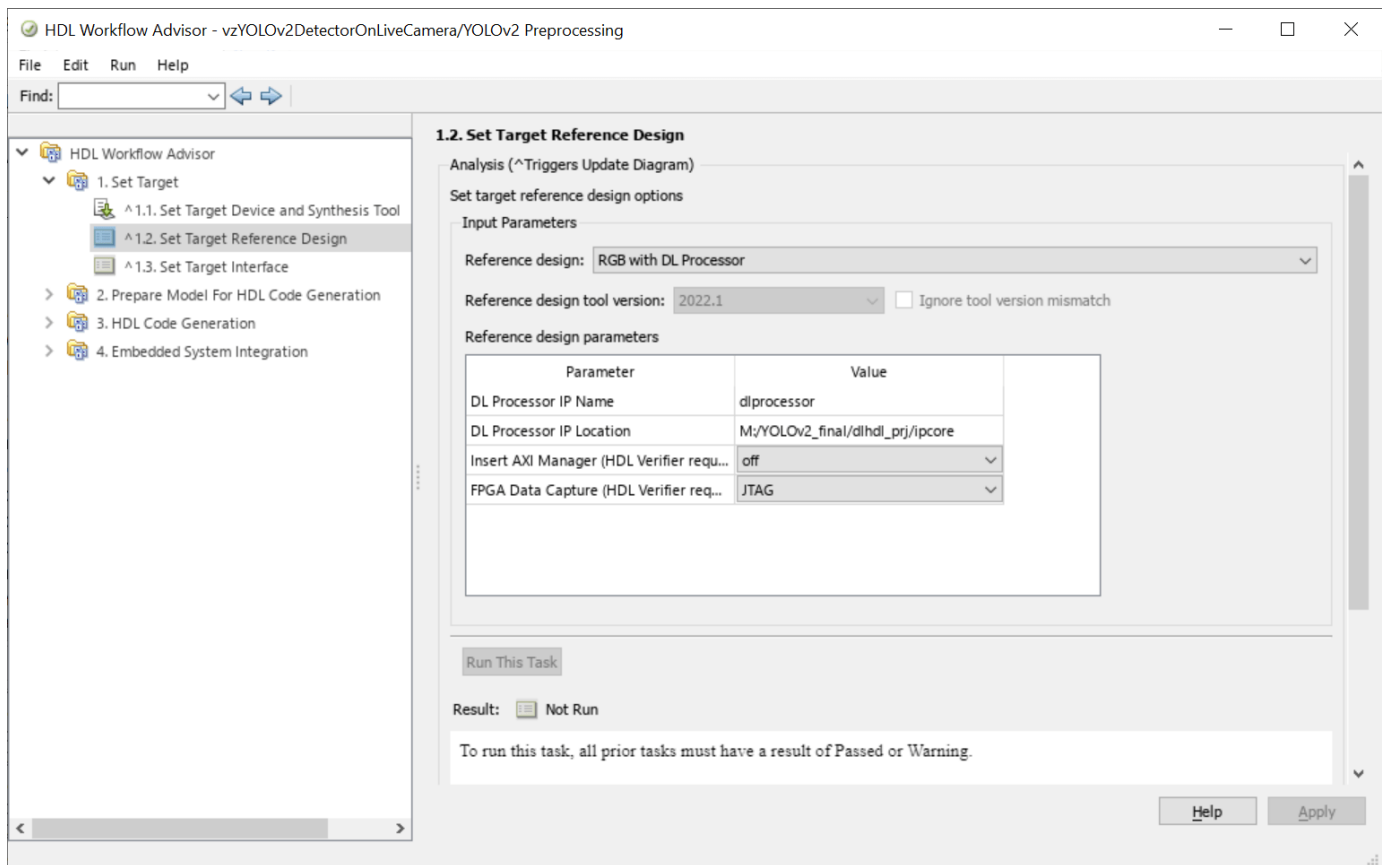
In step 1.3, map the target platform interfaces to the input and output ports of the DUT. The Pixel Streaming Data signals R,G,B of the algorithm will be mapped to the R,G and B signals of the Target Platform Interface. Similarly, the Pixel Control bus will be mapped to the Pixel Control bus signal in the Target Platform Interface. AXI4-Lite Interface: The DUTProcstart register is mapped to the AXI4-Lite register. When this register is written, it triggers the process of input handshaking logic. Choosing the AXI4-Lite interface directs HDL Coder to generate a memory-mapped register in the FPGA fabric. You can access this register from software running on the ARM processor. AXI4 Master DDR interface: The AXIWriteCtrlInDDR, AXIReadCtrlInDDR, AXIReadDataDDR, AXIWriteCtrlOutDDR, AXIWriteDataDDR and AXIReadCtrlOutDDR ports of DUT are mapped to AXI4 Master DDR interface. The Read Channel of the AXI4 Master DDR interface is mapped to the AXI4 Master DDR Read interface, and the Write Channel of the AXI4 Master DDR interface is mapped to the AXI4 Master DDR Write interface. This interface is used for the data transfer between the Preprocess DUT and the PL DDR. Using the Write Channel of this interface, the preprocessed data is written to the PL DDR which can then be accessed by the Deep Learning Processor IP.

AXI4 Master DL interface: The AXIReadDataDL, AXIReadCtrlInDL, AXIWriteCtrlInDL, AXIReadCtrlOutDL, AXIWriteDataDL and AXIWriteCtrlOutDL ports of DUT are mapped to AXI4 Master DL interface. The Read Channel of the AXI4 Master DL interface is mapped to the AXI4 Master DL Read interface, and the Write Channel of the AXI4 Master DL interface is mapped to the AXI4 Master DL Write interface. This interface is used for the communication between Preprocess DUT and the Deep Learning Processor IP. In this example, this interface is used for implementing input handshaking logic with Deep Learning Processor.

Port Name	Port Type	Data Type	Target Platform Interfaces	Interface Mapping
RIn	Inport	uint8	R Input [0:7] ⌄	[0:7]
GIn	Inport	uint8	G Input [0:7] ⌄	[0:7]
BIn	Inport	uint8	B Input [0:7] ⌄	[0:7]
CtrlIn	Inport	bus	Pixel Control Bus Input ⌄	
DUTProcstart	Inport	boolean	AXI4-Lite ⌄	x"100"
inputAddrDL	Inport	uint32	AXI4-Lite ⌄	x"104"
inputCounterReset	Inport	uint32	AXI4-Lite ⌄	x"10C"
AXIWriteCtrlInDDR	Inport	bus	AXI4 Master DDR Write ⌄	Write Slave to Master Bus ⌄
AXIReadCtrlInDDR	Inport	bus	AXI4 Master DDR Read ⌄	Read Slave to Master Bus ⌄
AXIReadDataDDR	Inport	ufix128	AXI4 Master DDR Read ⌄	Data ⌄
AXIReadDataDL	Inport	uint32	AXI4 Master DL Read ⌄	Data ⌄
AXIReadCtrlInDL	Inport	bus	AXI4 Master DL Read ⌄	Read Slave to Master Bus ⌄
AXIWriteCtrlInDL	Inport	bus	AXI4 Master DL Write ⌄	Write Slave to Master Bus ⌄
ROut	Outport	uint8	R Output [0:7] ⌄	[0:7]
GOut	Outport	uint8	G Output [0:7] ⌄	[0:7]
BOut	Outport	uint8	B Output [0:7] ⌄	[0:7]
CtrlOut	Outport	bus	Pixel Control Bus Output ⌄	
AXIReadCtrlOutDL	Outport	bus	AXI4 Master DL Read ⌄	Read Master to Slave Bus ⌄
AXIWriteDataDL	Outport	uint32	AXI4 Master DL Write ⌄	Data ⌄
AXIWriteCtrlOutDL	Outport	bus	AXI4 Master DL Write ⌄	Write Master to Slave Bus ⌄
AXIWriteCtrlOutDDR	Outport	bus	AXI4 Master DDR Write ⌄	Write Master to Slave Bus ⌄
AXIWriteDataDDR	Outport	ufix128	AXI4 Master DDR Write ⌄	Data ⌄
AXIReadCtrlOutDDR	Outport	bus	AXI4 Master DDR Read ⌄	Read Master to Slave Bus ⌄
ROIFramecount	Outport	uint32	AXI4-Lite ⌄	x"108"
FrameCount	Outport	uint32	AXI4-Lite ⌄	x"110"

Step 2 prepares the design for generation by doing some design checks.

Step 3 generates HDL code for the IP core.

Step 4.1 integrates the newly generated IP core into the larger Vision Zynq reference design.

In Step 4.2, the workflow generates a targeted hardware interface model and, if the Embedded Coder Zynq support package has been installed, a Zynq software interface model. Since this example uses the shipping example model, uncheck **Generate Simulink software interface model** and **Generate host interface script**.
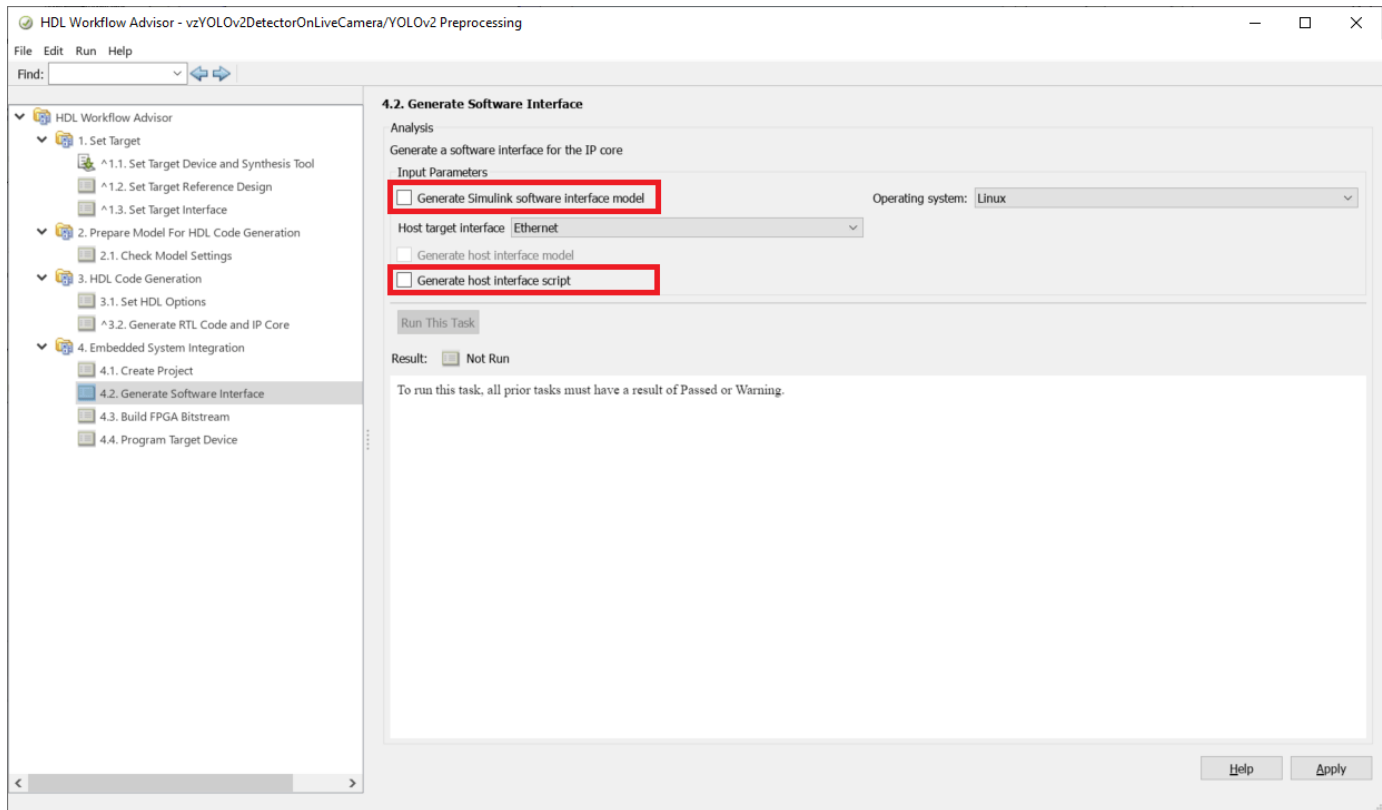
Click Run this task button with these settings. The rest of the workflow generates a bitstream for the FPGA, downloads it to the target, and reboots the board.

Because this process can take 3-4 hours, you can choose to bypass this step by using a pre-generated bitstream for this example that ships with product and was placed on the SDCard during setup.

Note: This bitstream was generated with the HDMI pixel clock constrained to 148.5 MHz for a maximum resolution of 1080p HDTV at 60 frames-per-second.

To use this pre-generated bitstream execute the following commands to copy the device tree file to the current working directory and to load the bitstream on hardware.

copyfile(fullfile(matlabshared.supportpkg.getSupportPackageRoot, "toolbox","shared","supportpackages","visionzynq","bin", "target","sdcard","visionzynq-zcu102-hdmicam","visionzynq-customtgt", "visionzynq-zcu102-hdmicam-dl.dtb"),"visionzynq-zcu102-hdmicam-dl.dtb");

vz = visionzynq();

changeFPGAImage(vz,'visionzynq-zcu102-hdmicam-dl-yolov2.bit', 'visionzynq-zcu102-hdmicam-dl.dtb')

To configure the Zynq device with this bitstream file at a later stage, execute the following commands:

To copy the dtb file to the current working directory, use this command

copyfile(fullfile(matlabshared.supportpkg.getSupportPackageRoot,
"toolbox","shared","supportpackages","visionzynq","bin", "target","sdcard","visionzynq-zcu102-
hdmicam","visionzynq-refdes", "visionzynq-zcu102-hdmicam-dl.dtb"),"visionzynq-zcu102-hdmicam-
dl.dtb");

vz = visionzynq();

downloadImage(vz,'FPGAImage', '<PROJECT_FOLDER>\vivado_ip_prj\vivado_prj.runs
\impl_1\design_1_wrapper.bit',' DTBImage', 'visionzynq-zcu102-hdmicam-dl.dtb')

Compile and Deploy YOLO v2 Deep Learning Network

Now that the bitstream is loaded in the above step, you can deploy the end to end DL Application on
the FPGA. Update the bitstream build information in the MAT file generated during the IP core
generation. The name of the MAT file is dlprocessor.mat and is located in cwd\dlhdl_prj\, where cwd is
your current working folder. Copy the file to the present working folder. This MAT file is generated
using the target platfom Generic Deep Learning Processor does not contain the Board/Vendor
information. Use updateBitstreamBuildInfo.m function to update the Board/Vendor information and
generate a new MAT file with the same name as generated bitstream.

bitstreamName = 'design_1_wrapper';

updateBitstreamBuildInfo('dlprocessor.mat', [bitstreamName,'.mat']);

Create a target object to connect your target device to the host computer.

hTarget = dlhdl.Target('Xilinx', 'Interface', 'Ethernet', 'IpAddr', '192.168.4.2');

Create a deep learning HDL workflow object using the dlhdl.Workflow class. Before running this
command, make sure that the generated bit file is available in the current working directory with the
name 'visionzynq-zcu102-hdmicam-dl-yolov2.bit'

hW = dlhdl.Workflow('Network', net, 'Bitstream', 'visionzynq-zcu102-hdmicam-dl-yolov2.bit', 'Target',
hTarget);

Compile the network, net using the dlhdl.Workflow object.

frameBufferCount = 2;

compile(hW, 'InputFrameNumberLimit', frameBufferCount);

Run the deploy function of the dlhdl.Workflow object to download the network weights and biases on
the Zynq UltraScale+ MPSoC ZCU102 board.

deploy(hW, 'ProgramBitStream', false);

Clear the DLHDL workflow object and hardware target.
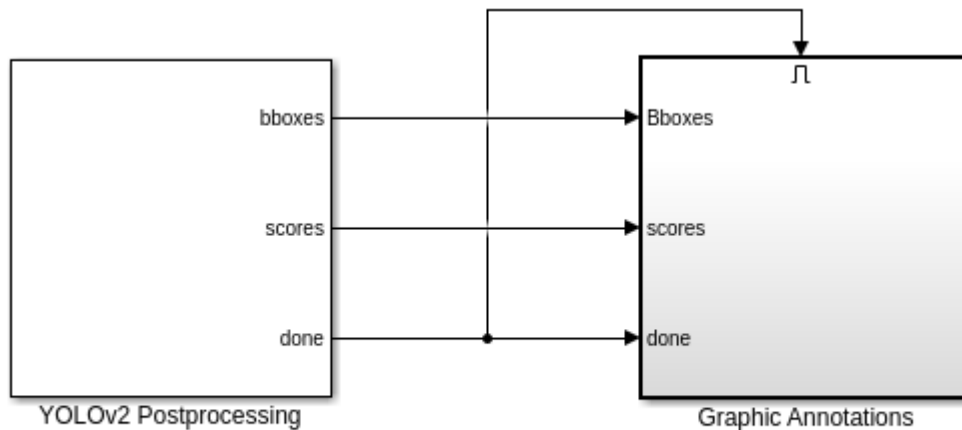
clear hW;

clear hTarget;

Software interface model

You can run this model in External mode on the ARM processor, or you can use this model to fully
deploy a software design. (This model can be deployed only if Embedded Coder and the Embedded
Coder Support Package for Xilinx Zynq Platform are installed.)

```
open_system('vzYOLOv2PostProcess');
```


Copyright 2022-23 The MathWorks, Inc.

Before running this model, you must perform additional setup steps to configure the Xilinx cross-compiling tools. For more information, see Setup for ARM Targeting. In the Postprocessing model, the YOLOv2 Postprocessing subsystem is same as the "Integrate YOLO v2 Vehicle Detector System on SoC" on page 3-41 The postprocessing model configures the DL processor for streaming mode up to a specified number of frames. The output data written to the PL DDR by DL Processor is read using AXI4 Stream IIO Read block.

Once the bounding boxes and scores are calculated in the YOLOv2PostprocessDUT block, the valid signal will be high. This valid signal goes to both draw Rect and set ROI blocks and are used for the synchronization between the input Frame written to the DDR and the bounding boxes and scores calculated. AXI4-Lite registers transfer the control signals between the FPGA and the ARM.

In this example, the software interface model contains only the postprocessing logic, and does not include a Video Capture HDMI block. This model is intended to run on the board independently from Simulink and does not return any data from the board. To view the output video in Simulink, you can use a different model that contains a Video Capture HDMI block and runs while your deep learning design is deployed and running on the board.

Open the 'vzYOLOv2PostProcess' model and click on 'Build, Deploy and Start' This mode runs the algorithm on the ARM processor on the Zynq board.

After opening the vzGettingStarted model, In Video Capture HDMI block, change the 'Video source' to 'HDMI input', 'Frame size' to '1080p HDTV (1920x1080p)', 'Pixel Format' to 'RGB' and 'Capture Point' to 'Output from FPGA user logic (B)' and in To Video Display block, change the 'Input Color Format' to 'RGB' and run the model. The bounding boxes and scores that are calculated in the ARM are overlaid on the corresponding frame and are displayed by the 'To Video Display' block in the vzGettingStarted model.

To stop the executable on ARM, run this command:

```
vz.stopExecutable('/tmp/vzYOLOv2PostProcess.elf');
```

See Also

Related Examples

- "Deploy and Verify YOLO v2 Vehicle Detector on FPGA" on page 3-14
- "Integrate YOLO v2 Vehicle Detector System on SoC" on page 3-41

More About

- "Deep Learning Processing of Live Video" (Vision HDL Toolbox Support Package for Xilinx Zynq-Based Hardware)

Vertical Video Flipping Using External Memory

This example shows how to design an application to flip an incoming video stream vertically.

You will use external memory to store a video frame to accomplish this task. You can use this technique for designing and implementing other vision applications requiring access to the external memory.

Supported hardware platform

- Xilinx® Zynq® ZC706 evaluation kit + FMC-HDMI-CAM mezzanine card

Introduction

To flip the incoming video stream from HDMI source (RX), the FPGA logic writes the video frames a line at a time to the external memory. Later, the FPGA logic reads the stored image back, a line at a time in the reverse order, thus flipping the image vertically. The read video frame is then sent to the HDMI Out (TX). Video frames are stored in a ping-pong buffer in the PL-DDR, which enables independent memory write and read operations. A separate, PS-DDR is used for storing video frames during transfer of data to the HDMI Out. The diagram below highlights the overall data flow.
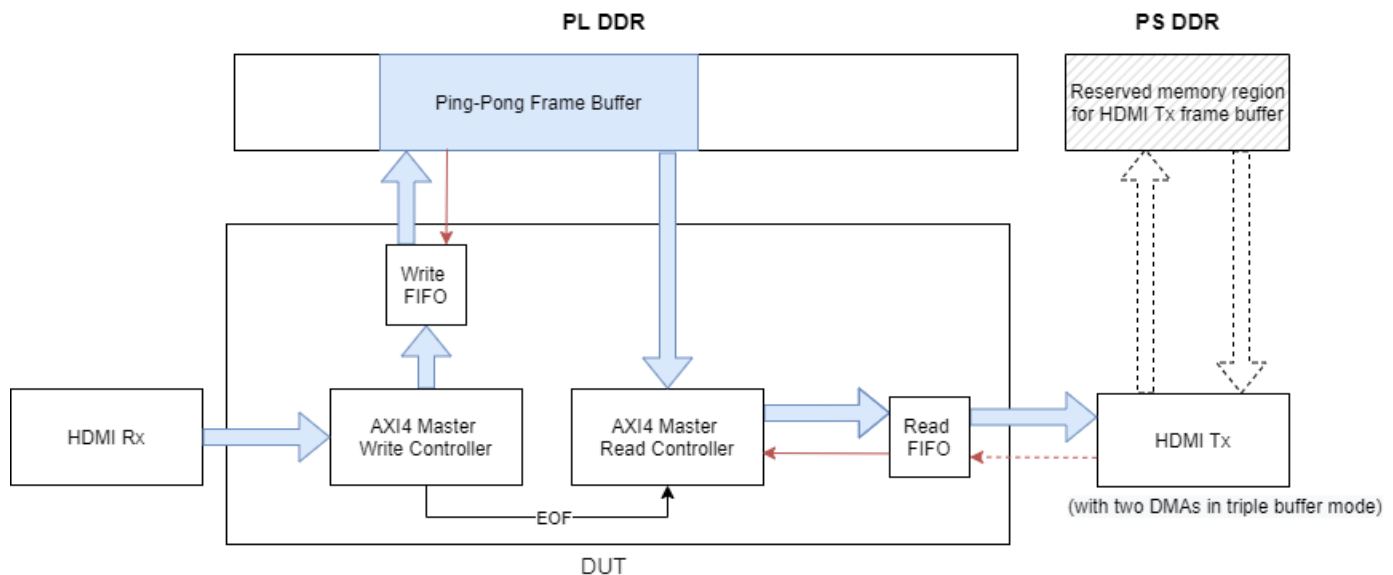
Modeling

In the top model, soc_video_flipping_top the FPGA logic is connected with the external memory and the HDMI In/Out blocks.

Vertical Video Flipping Using External Memory
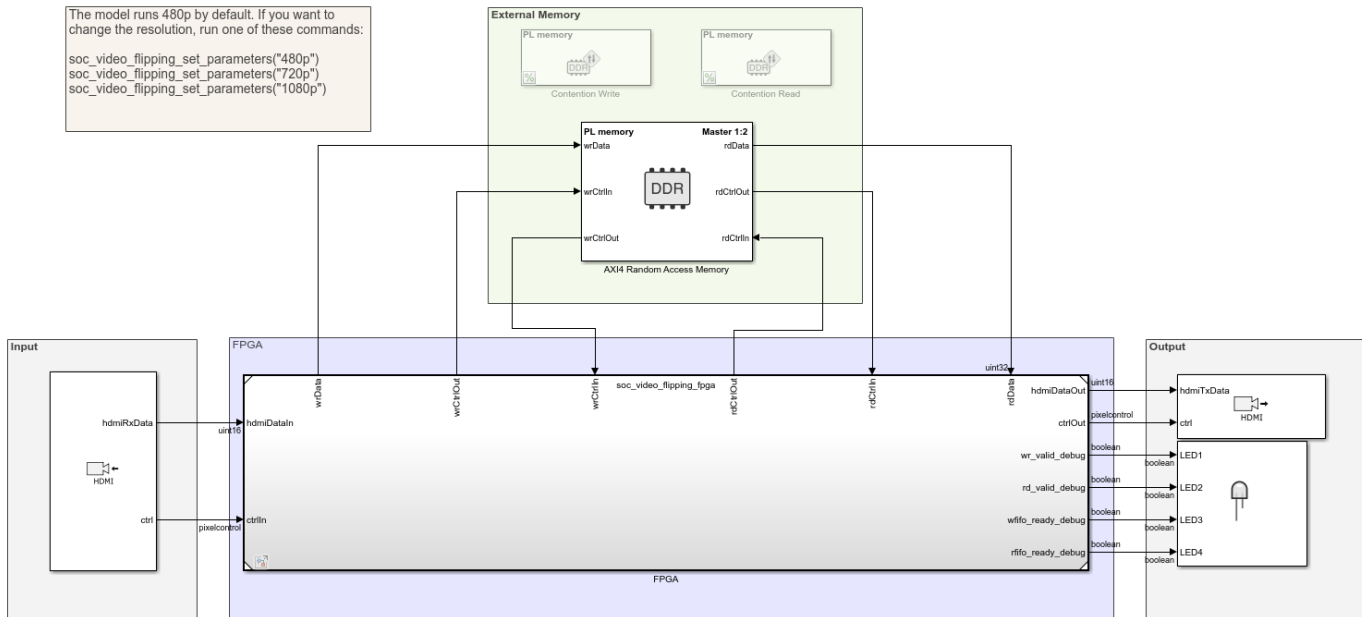
The soc_video_flipping_fpga model includes the FPGA logic. It is linked as a referenced model from the top model. This image shows the contents of the VideoFlipping subsystem inside the soc_video_flipping_fpga model.
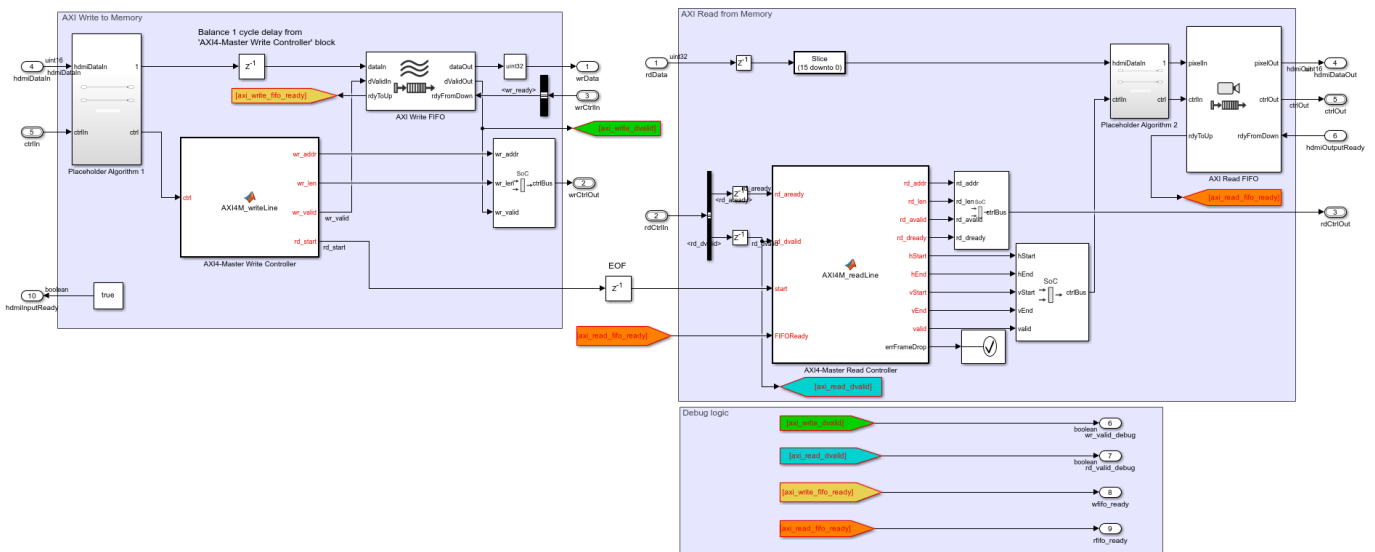
The FPGA logic consists of four key components:

- **AXI4 Master Write Controller** receives video from HDMI Rx and writes the data into DDR. One line of data is written per burst. When one frame finishes, it sends the end of frame (EOF) signal to trigger AXI4 Master Read Controller.

- **AXI4 Master Read Controller** reads the lines of the video frame from the external memory ping-pong buffer in reverse order. One line of data is read per burst. New read requests are paused if the downstream Read FIFO block is not ready to process data. When a frame is fully read from memory, the Read Controller waits for the next EOF signal from the Write Controller to start reading in a new frame. If the memory controller doesn't have enough bandwidth, the read controller may be still processing the earlier frame when the write controller finishes writing in the next frame. In this case, the Read Controller will throw an error using errFrameDrop signal.

- **Write FIFO** buffers the data written to the ping-pong buffer when the DDR memory controller asserts the backpressure signal (highlighted red arrow), to allow video data from HDMI source to be processed. The Write FIFO should be large enough to prevent overflow and accommodate any delay in writing to the external memory. In this example the depth of the Write FIFO is set to 2048 to accommodate 1 HD line of backpressure.

- **Read FIFO** buffers the data from the Read Controller when the DMA in the HDMI Tx asserts the backpressure signal (highlighted red arrow). The backpressure is propagated to the upstream AXI4 Master Read Controller to stop requesting data from DDR. Notice that since the AXI4 Master Read Controller will not pause during the read burst, it is important to make sure that the Read FIFO has enough room to store data even after its ready signal de-asserts. In this example, the depth of the Read FIFO is set to 2048 and the almost full threshold is set to 128. When the FIFO has 128 samples, the Read FIFO sends 'Full' signal to upstream block to stop any new read requests. In the meanwhile, it can buffer 2048-128 = 1920 samples without an overflow. The setting is sufficient even for a 1080p frame.

In the hardware implementation, the HDMI Tx includes two DMA frame buffers and one video timing controller (VTC), for robust and tear free video output. The DMAs may send backpressure to DUT when the memory controller is busy with other read or write transactions. In the FPGA model, the backpressure signal `hdmiOutputReady` is set to always true for simulation only (which indicates that the memory controller is always available). In practice, this signal often toggles between high and low. The Read FIFO block in the DUT is used to handle this backpressure.

Simulation

The memory bandwidth requirement must be considered when designing an application that interfaces with external memory. What is the rate that you need to transfer data to/from memory to satisfy the requirements of your algorithm? Specifically, for vision applications, what is the frame-size and frame-rate that you must be able to maintain?

For the selected ZC706 board, PL DDR controller is configured with 64-bit AXI4-Slave interface running at 200 MHz. The resulting bandwidth is 1600 MB/s. Let's first evaluate if the memory bandwidth is sufficient to maintain a 1920x1080p video stream at 60 frames-per-second. As the video format is YCbCr 4:2:2, we require 2 bytes-per-pixel. However, for the DUT AXI4 read and write, each pixel is zero-padded to 4 bytes, this equates to a throughput requirement of

$$2 \times 4 \times 1920 \times 1080 \times 60 = 995.328 \, \text{MB/s}$$

The calculated throughput satisfies the bandwidth requirement.

To simulate 1080p 60fps case, run the following command and then simulate the model

```
soc_video_flipping_set_parameters("1080p")
```

The output is shown as below.

If you want to model another DUT accessing the same external memory, you can use memory traffic generator block to simulate the contention before the implementation, so you can save effort on hardware debugging. In this model, two memory traffic generator blocks, `Contention Write` and `Contention Read` block, are modeled to mimic the AXI transactions of another frame buffer. The throughput of two memory traffic generators is calculated as

$$2 \times 2 \times 1920 \times 1080 \times 60 = 497.664 \, \text{MB/s}$$

The total required bandwidth of memory controller is $497.664 + 995.328 = 1492.992 \, \text{MB/s}$ Which is less than its maximum bandwidth 1600 MB/s. Uncomment Contention Write and Contention Read block, simulate the model for 1080P again, the output is teared at the bottom. You will also get assertion by 'soc_video_flipping_fpga/VideoFlipping/Assertion' block that frame get dropped.
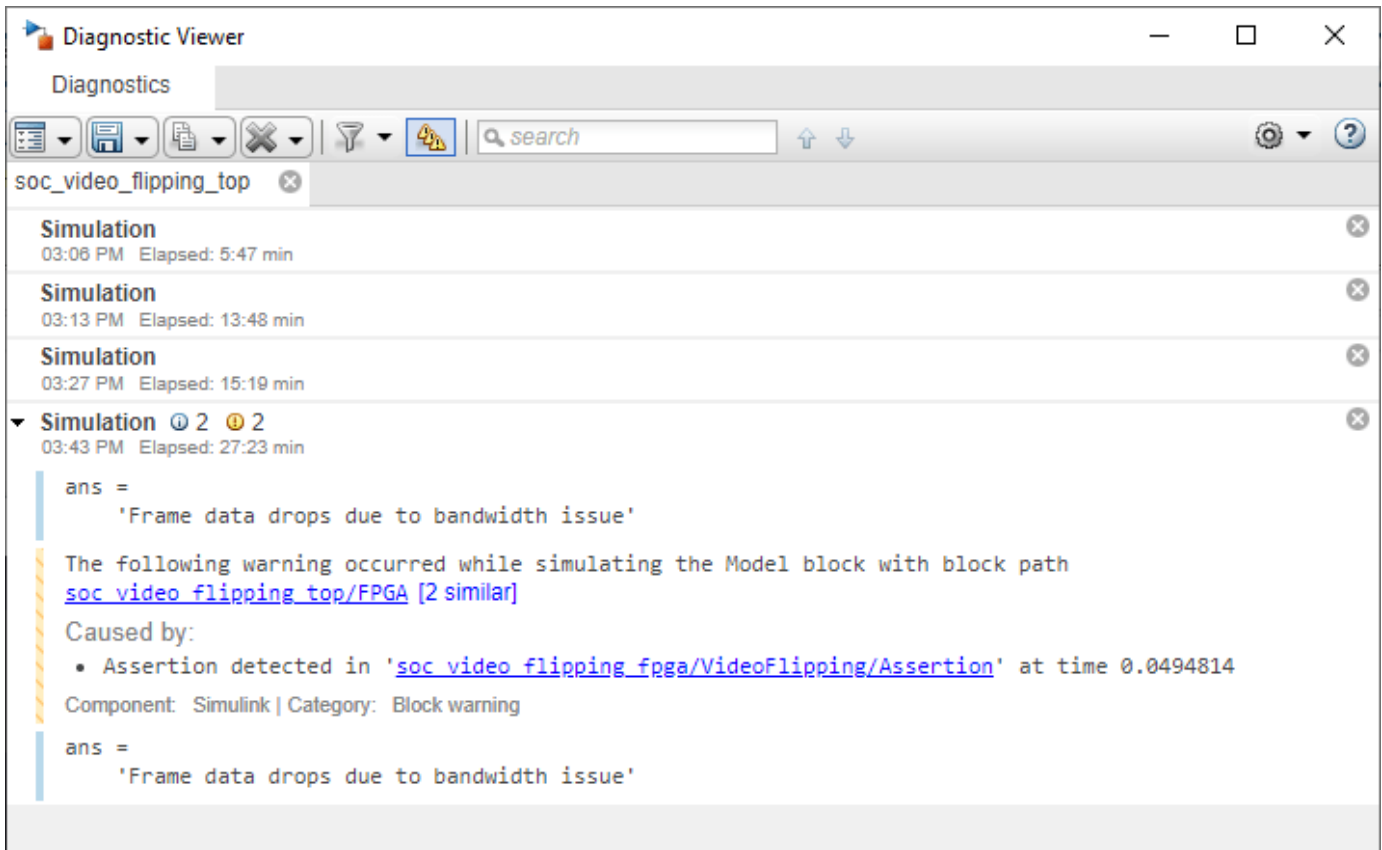
Simulation shows that memory controller can not meet the bandwidth requirement for the DUT and memory traffic generators. You may need to consider reducing the frame resolution, or making DUT algorithm more efficient, for example implementing pack/unpack 16bit YCbCr pixel to 32bit, other than zero padding.

The top model runs in Accelerator mode by default. If you want to inspect the logged data in Logic Analyzer, change the top model to Normal mode. It is best to simulate with 480p or smaller frame size for faster results.
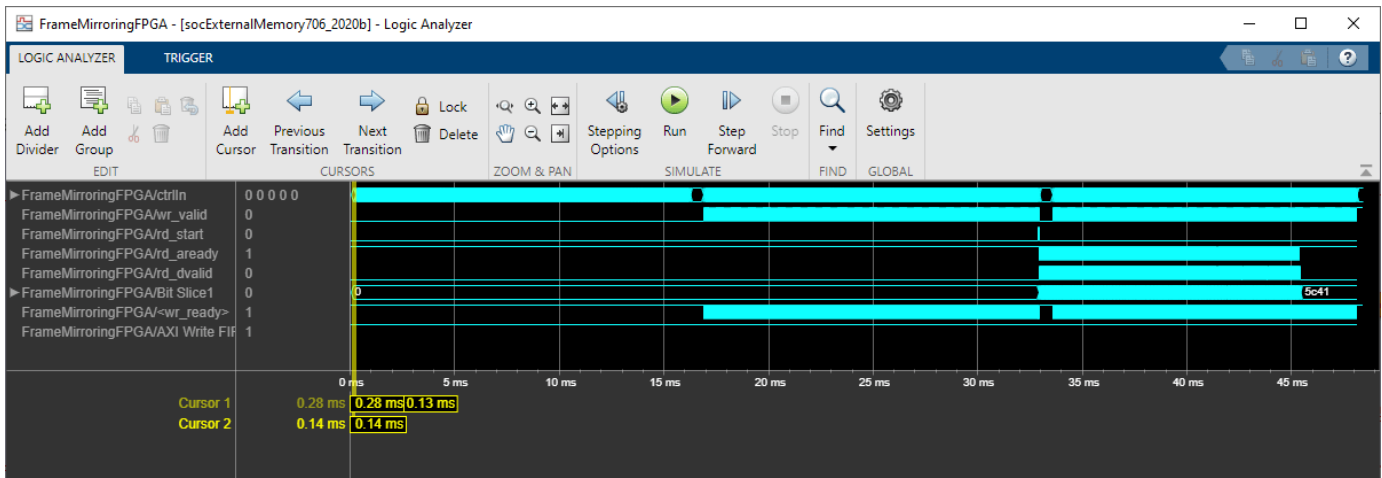
Implementation

Following products are required for this section:

- HDL Coder™

Before implementation,

- Set the simulation mode to 'Normal' on the top model.
- Comment out `Contention Write` and `Contention Read` blocks.
- Run this command if the model has been changed for the frame size other than 1080p.

```
soc_video_flipping_set_parameters("1080p")
```

To implement the model on a supported SoC board use the SoC Builder (SoC Blockset) tool. Make sure you have installed required products and FPGA vendor software before implementation. To open **SoC Builder** click, **Configure, Build, & Deploy** button in the toolstrip and follow these steps:

1. On the **Setup** screen, select **Build model**. Click **Next**.
2. On the **Select Build Action** screen, select **Build, load and run**. Click **Next**.
3. On the **Select Project Folder** screen, specify the project folder. Click **Next**.
4. On the **Review Memory Map** screen, to view the memory map, click **View/Edit**. Click **Next**.
5. On the **Validate Model** screen, to check the compatibility of the model for implementation, click **Validate**. Click **Next**.
6. On the **Build Model** screen, to build the model, click **Build**. An external shell opens when FPGA synthesis begins. Click **Next**.
7. On the **Connect Hardware** screen, to test the connectivity of the host computer with the SoC board, click **Test Connection**. To go to the **Run Application** screen, click **Next**.

The FPGA synthesis often takes more than 30 minutes to complete. To save time, you can use the provided pregenerated bitstream by following these steps.

1. Close the external shell to terminate synthesis.
2. Copy pregenerated bitstream to your project folder by running this `copyfile` command below.
3. Load the pregenerated bitstream and run the model on the SoC board by clicking **Load and Run**.

```
copyfile(fullfile(matlabshared.supportpkg.getSupportPackageRoot, ...
 'toolbox','soc','supportpackages','xilinxsoc','xilinxsocexamples', ...
 'bitstreams','soc_video_flipping_top-zc706.bit'),'./soc_prj');
```

Four LEDs on the ZC706 are driven by signals and can be used for debugging the design:

- GPIO_LED_LEFT is driven by AXI4 Master write data valid. It should be on or blinking when the application is running.
- GPIO_LED_CENTER is driven by AXI4 Master read data valid. It should be on or blinking when the application is running.
- GPIO_LED_RIGHT is driven by Write FIFO ready. It should be always on, otherwise AXI4 Write data get dropped.
- GPIO_LED_0 is driven by Read FIFO ready. It could be on or off. The data won't get dropped in this FIFO because the upstream controller will handle this backpressure properly.

Conclusion

This example shows modeling of AXI4 Master interfaces for accessing external memory in random fashion using SoC Blockset. You can use this technique to model vision applications involving external memory. One such example is "Contrast Limited Adaptive Histogram Equalization with External Memory" on page 3-89 which builds further on this example.

See Also

Memory Channel | Memory Controller | Memory Traffic Generator

Related Examples

- "Contrast Limited Adaptive Histogram Equalization with External Memory" on page 3-89

Rotate Image by Small Acute Angle

This example shows how to implement an image rotation algorithm for small acute angles for FPGA.

Image rotation is the movement of an image around a fixed point. It is one of the most common affine transforms and is fundamental to many computer vision applications like feature extraction and matching. This equation represents the affine transform that rotates new coordinates (x, y) from original coordinates (u, v) by rotation angle θ.

$$\begin{bmatrix} x \\ y \end{bmatrix} = \begin{bmatrix} cos(\theta) & -sin(\theta) \\ sin(\theta) & cos(\theta) \end{bmatrix} \begin{bmatrix} u \\ v \end{bmatrix}$$

This implementation is based on the `imrotate` (Image Processing Toolbox) function.

This example computes the transformation matrix for an angle in the range (–10, 0) and (0, 10) by using the `ComputeSmallAngleAffineTransform.m` function. The transformation matrix returned by this function is an input to the hardware algorithm. The hardware algorithm performs an affine transform and calculates the output pixel intensities by using bilinear interpolation. This implementation does not require external DDR memory and instead uses the on-chip block RAM to store and resample the output pixel intensities.

Image Rotation Algorithm

The image rotation algorithm uses a reverse mapping technique to map the pixel locations of the output rotated image to the pixels in the input image. This diagram shows the different stages of the algorithm.
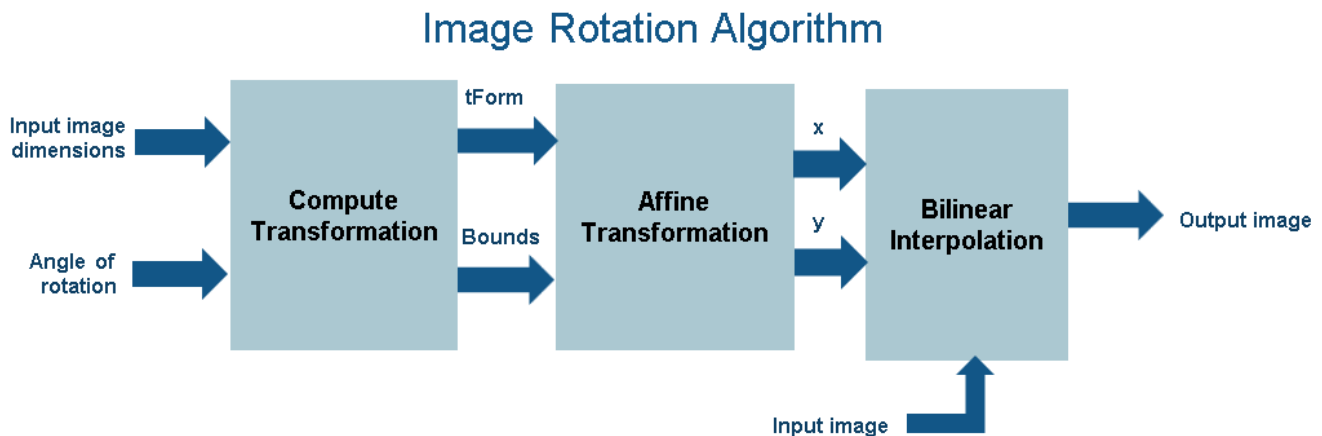
Compute Transformation This stage computes transformation parameters using the input image dimensions and the angle of rotation, θ. The transformation parameters that this stage outputs include the output bounds and the transformation matrix $tForm$. The bounds help compute the integer pixel coordinates of the output rotated image, and $tForm$ transforms the integer pixel coordinates in the output rotated image to corresponding coordinates of the input image.

Affine Transform An affine transform is a geometric transformation that translates a point in one image plane onto another image plane by preserving the images collinearity. Collinearity means that

all points on a line in the input image still form that line after transformation. Image rotation, maps integer pixel coordinates in the output rotated image to the corresponding coordinates of the input image by using the transformation matrix, $tForm$. If (u, v) is an integer pixel coordinate in the rotated output image and (x, y) is the corresponding coordinate of the input image, then this equation describes the transformation.

$$[x \; y \; z]_{1X3} = [u \; v \; 1]_{1X3} * tForm^{-1}_{3X3}$$

Bilinear Interpolation The rotation algorithm can produce (x, y) coordinates that are noninteger values. To generate the intensity of pixels at each integer position, a resampling technique like interpolation must be used. This example uses bilinear interpolation to resample the image intensity values corresponding to the generated coordinates.

In the equation and the diagram, (x, y) is the coordinate of the input pixel generated by the affine transform stage. $I1$, $I2$, $I3$, and $I4$ are the four neighboring pixels, and $deltaX$ and $deltaY$ are the displacements of the target pixel from its neighboring pixels. This stage of the algorithm computes the weighted average of the four neighboring pixels by using this equation.

$$outputPixel = I1(1 - deltaX)(1 - deltaY) + I2(deltaX)(1 - deltaY) + I3(1 - deltaX)(deltaY) + I4(deltaX)(deltaY)$$
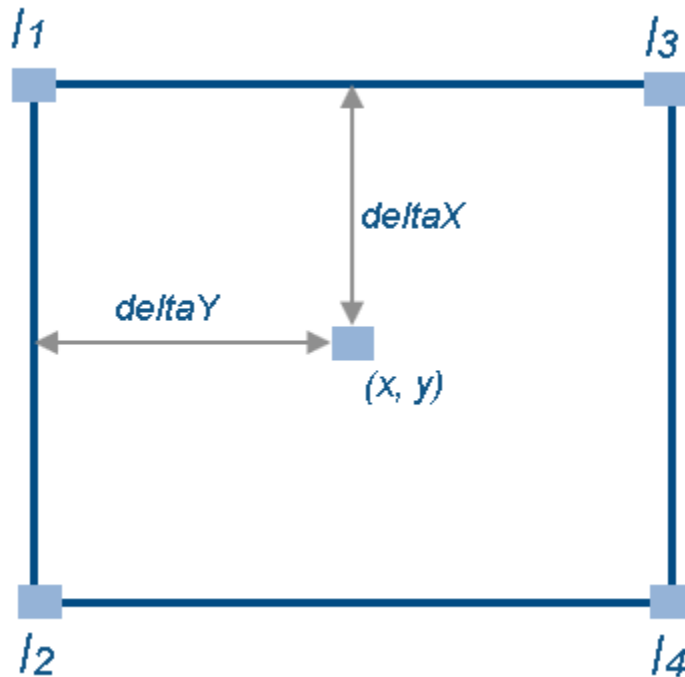
**HDL Implementation**

This figure shows the top-level view of the `ImageRotationHDL` model. The InputImage block imports the image from files. The Frame To Pixels block converts the input image frames to a pixel stream with a `pixelcontrol` bus for input to the `ImageRotationHDLALgorithm` subsystem. This subsystem rotates the input image by an angle that you can specify by using the mask of the

Transform block. The Pixels To Frame block converts the stream of output pixels back to frames. The `ImageViewer` subsystem displays the input frame and the corresponding rotated output.

```
open_system('ImageRotationHDL');
set(allchild(0),'Visible','off');
```
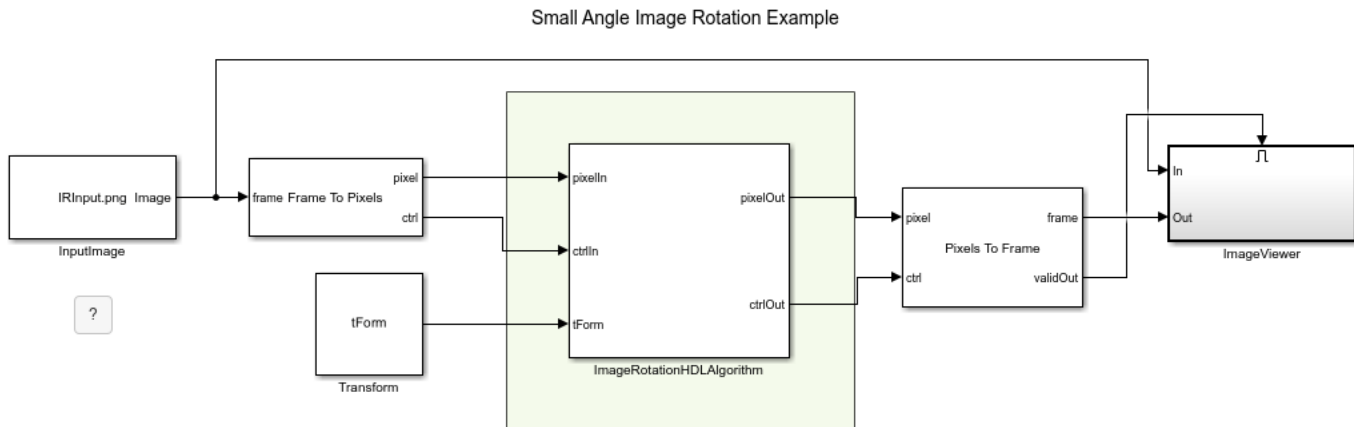

The `InitFcn` callback of the example model computes $tForm$ by calling the `ComputeSmallAngleAffineTransform.m` function. This function takes angle of rotation and input image dimensions as the input. You can set these values in the mask of the Transform block. Alternatively, you can generate your own transformation matrix (flattened to a 6-by-1 vector, because the last column of $tForm$ is redundant) and give it as an input to the `ImageRotationHDLAlgorithm` subsystem.

In the `ImageRotationHDLAlgorithm` subsystem, the `GenerateControl` subsystem generates a control signal `pixelcontrol` bus from the input **ctrl** bus depending upon the displacement parameter. The `CoordinateGeneration` subsystem generates the row and column pixel coordinates (u, v) of the output rotated image. It uses two HDL counters to generate the row and column coordinates. The `AffineTransform` subsystem maps these coordinates onto their corresponding row and column coordinates, (x, y), of the input image.

The `AddressGeneration` subsystem calculates the addresses of the four neighbors of (x, y) required for interpolation. This subsystem also computes the parameters $deltaX$, $deltaY$, $Bound$, and $indexVector$, which are used for bilinear interpolation.

The `Interpolation` subsystem stores the pixel intensities of the input image in a memory. To calculate each rotated output pixel intensity, the subsystem reads the four neighbor pixel values and computes their weighted sum.

```
open_system('ImageRotationHDL/ImageRotationHDLAlgorithm','force');
```
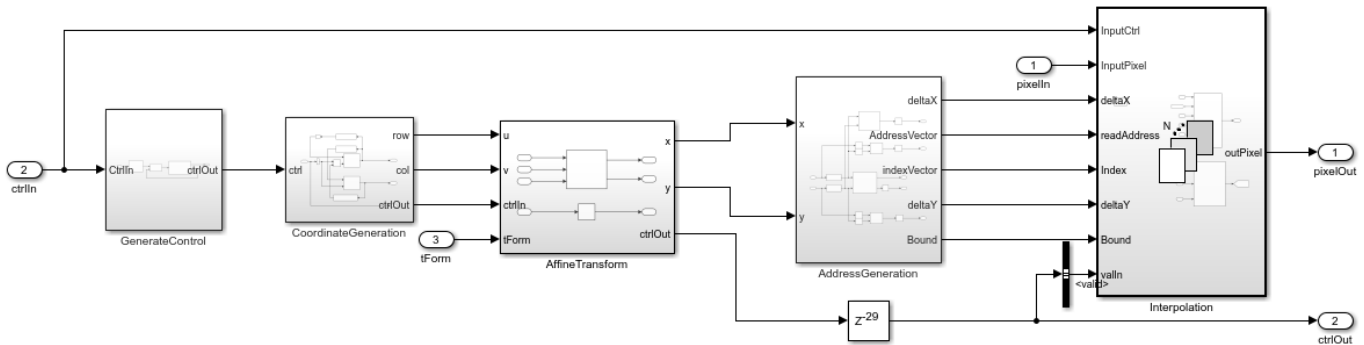
Affine Transformation

The HDL implementation of the affine transformation multiplies the coordinates *[u v 1]* with the transformation matrix, $tForm$ (flattened to a 6-by-1 vector, because the last column of $tForm$ is redundant). The `ComputeSmallAngleAffineTransform.m` function, called in the `InitFcn` callback of the model, generates the $tForm$ matrix. The `Transformation` subsystem implements the matrix multiplication with Product blocks which multiply the integer coordinates of output image by each element of the $tForm$ matrix. For this operation, the $tForm$ is split from a vector into individual elements by using a Demux block.

```
open_system('ImageRotationHDL/ImageRotationHDLAlgorithm/AffineTransform','force');
```
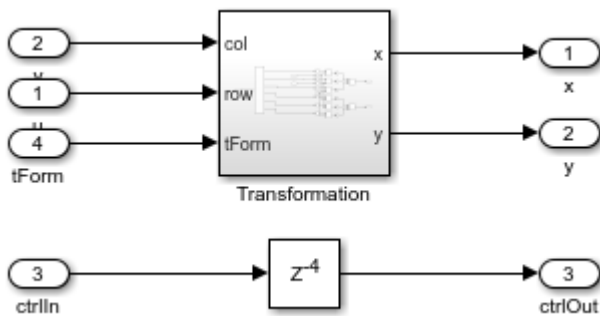

Address Generation

The `AddressGeneration` subsystem takes the mapped coordinate of the input raw image (x, y) as input and then calculates the displacement $deltaX$ and $deltaY$ of each pixel from its neighboring pixels. The subsystem also rounds the coordinates to the nearest integer toward negative infinity.

```
open_system('ImageRotationHDL/ImageRotationHDLAlgorithm/AddressGeneration','force');
```
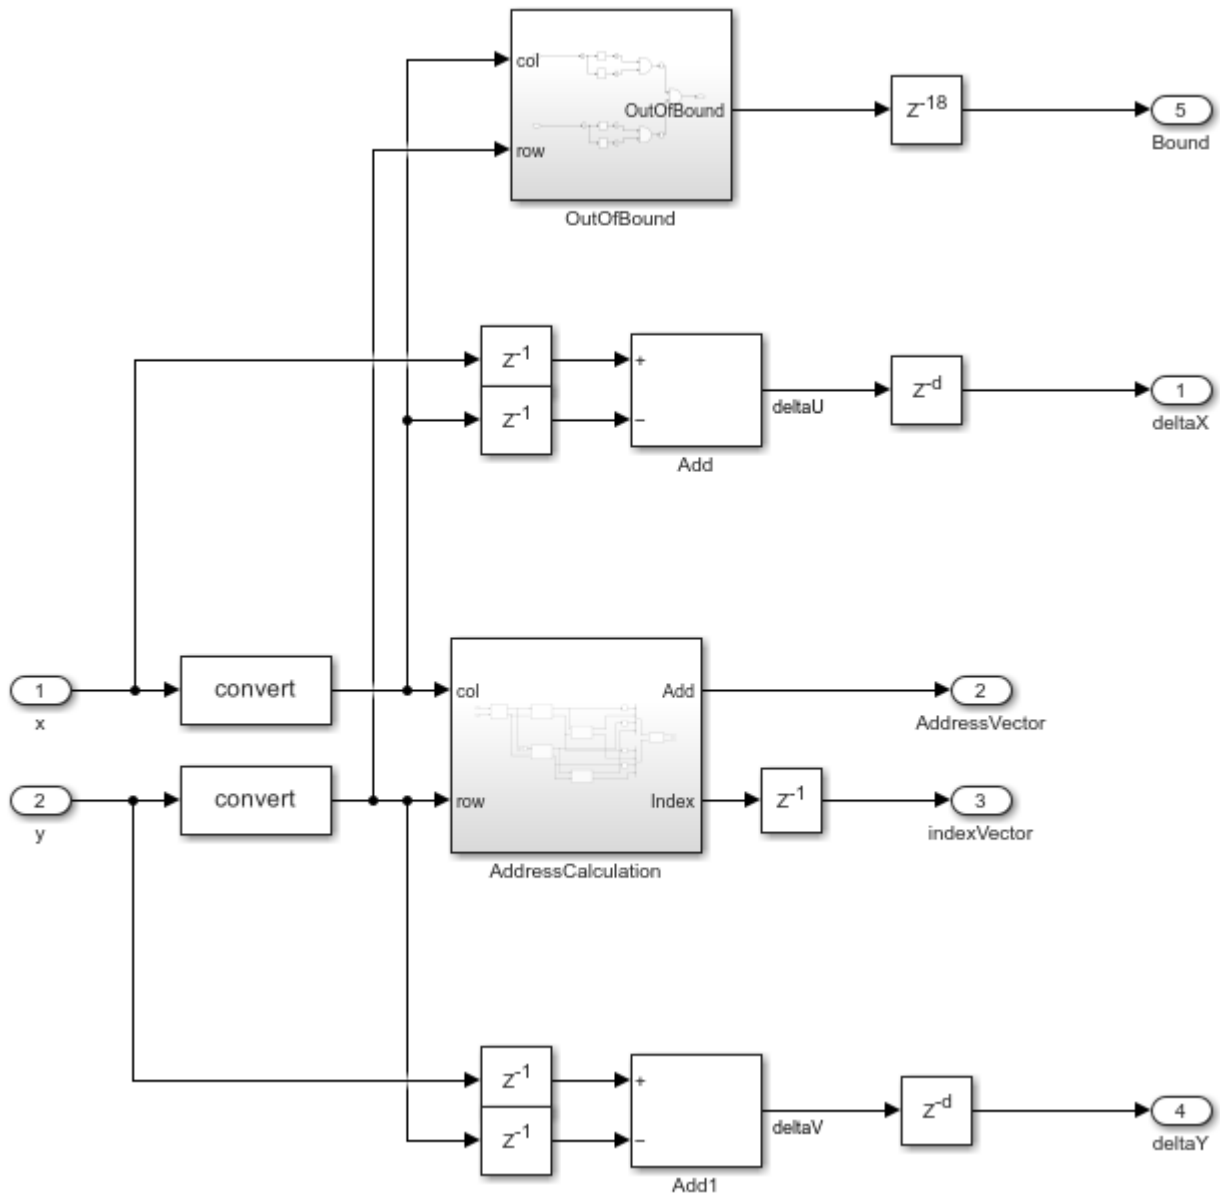
The `AddressCalculation` subsystem checks the coordinates against the bounds of the input images. If any coordinate is outside the image dimensions, that coordinate is capped to the boundary value for further processing. Next, the subsystem calculates the index of the address of each of the four neighborhood pixels in the CacheMemory block. The index represents the column of the cache. The subsystem finds the index for each address by using the even and odd nature of the incoming column and row coordinates, as determined by the Extract Bits block.

```
% ==========================
% |Row   || Col  || Index ||
% ==========================
% |Odd  || Odd  ||   1   ||
% |Even || Odd  ||   2   ||
% |Odd  || Even ||   3   ||
```

```
% |Even || Even ||   4   ||
% =========================
```

The address of the neighborhood pixels is generated using this equation.

$$Address = (\frac{SizeOfColumn}{2} * nR) + nC$$

nR is the row coordinate and nC is the column coordinate. When row is even, then $nR = \frac{row}{2} - 1$. When row is odd, then $nR = \frac{row-1}{2}$. When col is even, then $nC = \frac{col}{2}$. When col is odd, then $nC = \frac{col+1}{2}$.

The IndexChangeForMemoryAccess MATLAB Function block in the `AddressCalculation` subsystem rearranges the addresses in increasing order of their indices. This operation ensures the correct fetching of data from the CacheMemory block. The addresses are given as input to the CacheMemory block, and $index$, $deltaX$, and $deltaY$ are passed to the `Interpolation` subsystem.

The `OutOfBound` subsystem checks whether the (x, y) coordinates are out of bounds (that is, if any coordinate is outside the image dimensions). If the coordinate is out of bounds, the corresponding output pixel is set to an intensity value of 0.

After all of the addresses and their corresponding indices are generated, a Vector Concatenate block creates vectors of the addresses and indices.

Interpolation

The `Interpolation` subsystem is a For Each block which replicates its operation depending on the dimensions of the input pixel. For example, if the input is an RGB image, then the input pixel dimensions are 1-by-3, and the model includes 3 instances of this operation. Using the For Each block enables the model to support RGB input or grayscale input. The operation inside the For Each subsystem comprises two subsystems: `BilinearInterpolation` and `CacheMemory`.

```
open_system('ImageRotationHDL/ImageRotationHDLAlgorithm/Interpolation','force');
```
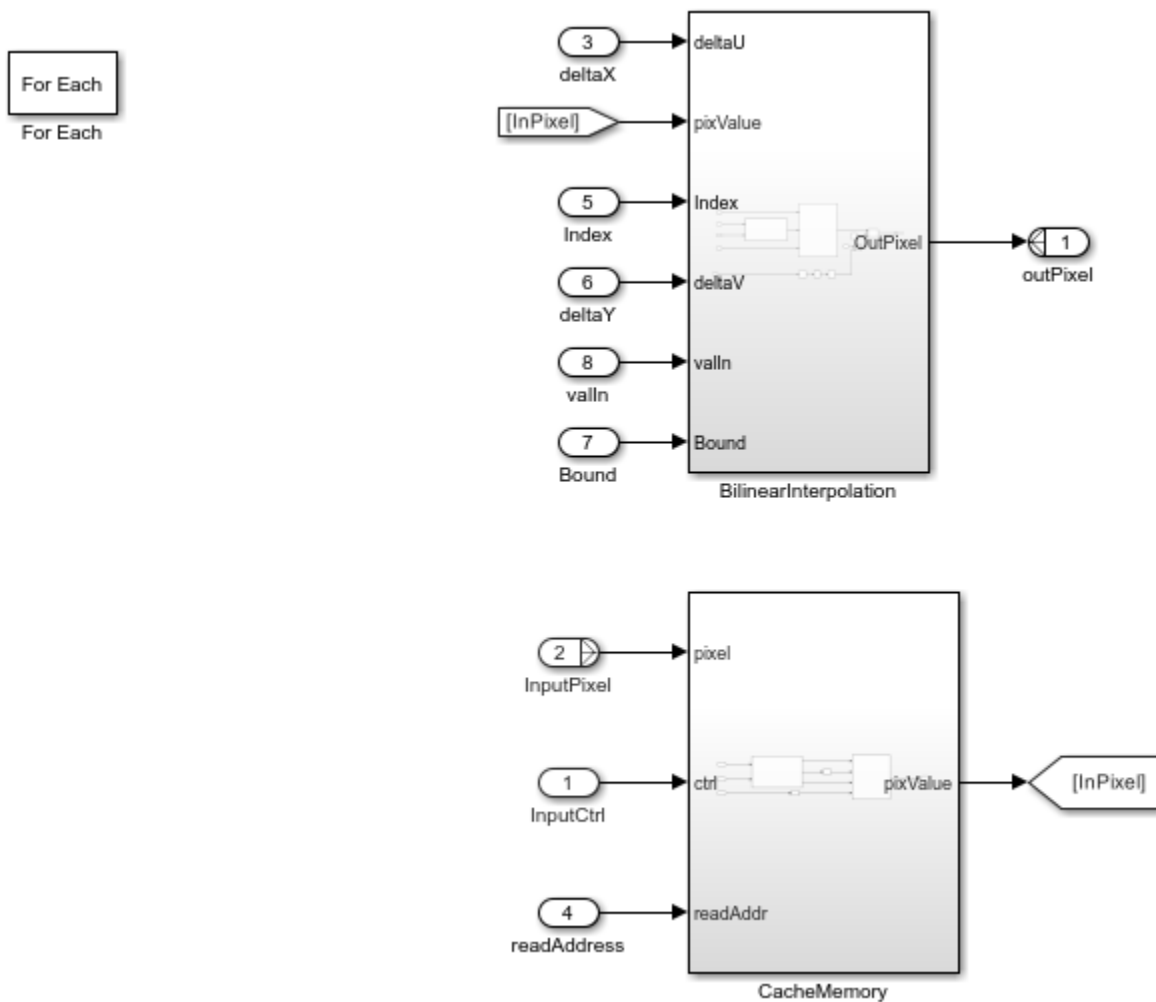
Cache Memory

The CacheMemory subsystem contains a Simple Dual Port RAM block. The subsystem buffers the input pixels to form [Line 1 Pixel 1 | Line 2 Pixel 1 | Line 1 Pixel 2 | Line 2 Pixel 2] in the RAM. This configuration enables the algorithm to read all four neighboring pixels in one cycle. The required size of the cache memory is calculated from the *offset* output of the ComputeSmallAngleAffineTransform.m function. The offset is the sum of maximum deviation and the first row map. The first row map is the maximum value of the input image row coordinate that corresponds to the first row of the output rotated image. The maximum deviation is the greatest difference between the maximum and minimum row coordinates for each row of the input image row map.

The WriteControl subsystem forms vectors of incoming pixels, write enables, and write addresses. The AddressGeneration subsystem provides a vector of read addresses. The vector of pixels that are read from the RAM becomes the input to the BilinearInterpolation subsystem.

```
open_system('ImageRotationHDL/ImageRotationHDLAlgorithm/Interpolation/CacheMemory','force');
```
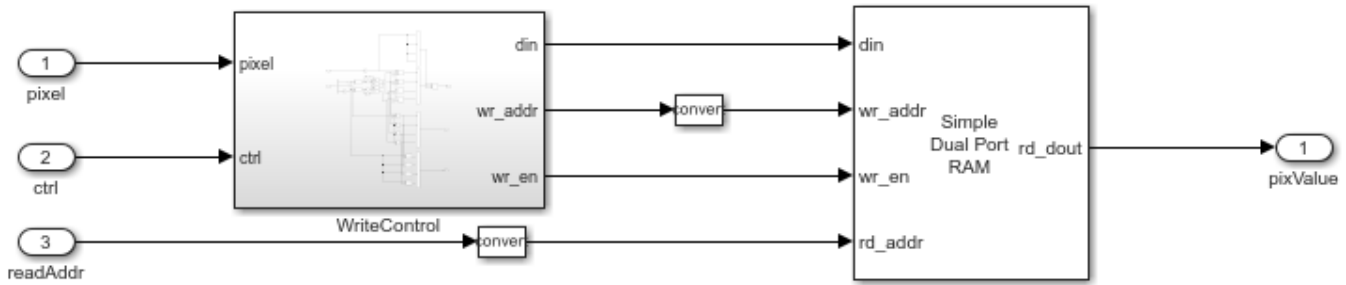
Bilinear Interpolation

The `BilinearInterpolation` subsystem rearranges the vector of read pixels from the cache to their original indices. Then, the `BilinearInterpolationEquation` subsystem calculates a weighted sum of the neighborhood pixels by using the bilinear interpolation equation mentioned in the Image Rotation Algorithm section. The result of the interpolation is the value of the output rotated pixel.

```
open_system('ImageRotationHDL/ImageRotationHDLAlgorithm/Interpolation/BilinearInterpolation','fo
```
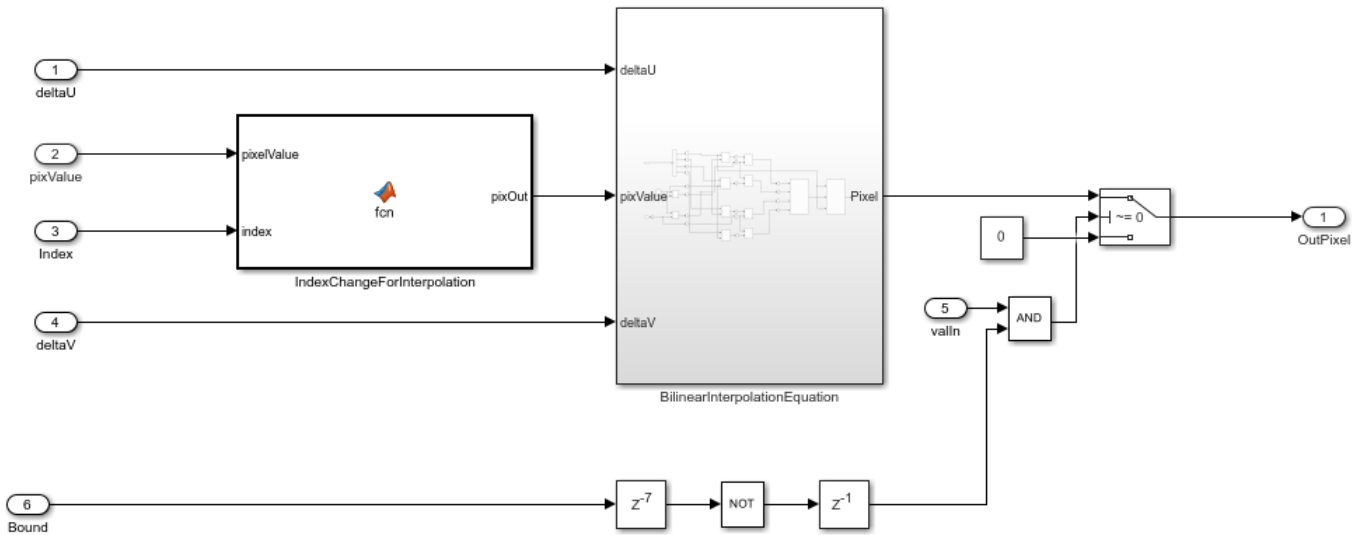

Simulation and Results

This example uses a 480p RGB input image. The input pixels use the `uint8` data type. The example supports either grayscale or RGB input images. This example supports acute angles in the range (–10, 0) and (0, 10). Angles greater than 10 degrees require much higher BRAM resources.

This figure shows the input image and the corresponding output image rotated by an angle of 7 degrees. The results of the `ImageRotationHDL` model for this input matches the output of the `imrotate` function.
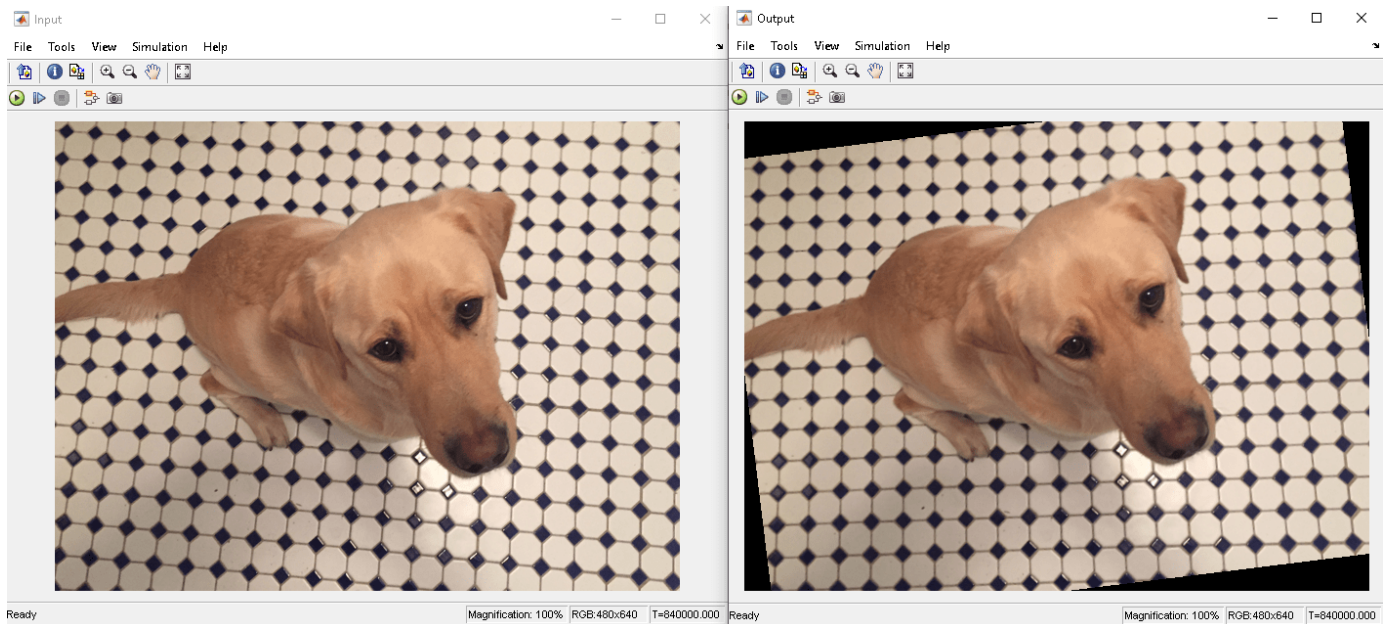
To check and generate the HDL code referenced in this example, you must have the HDL Coder™ product.

To generate the HDL code, enter this command.

```
makehdl('ImageRotationHDL/ImageRotationHDLAlgorithm')
```

To generate the test bench, enter this command.

```
makehdltb('ImageRotationHDL/ImageRotationHDLAlgorithm')
```

This design was synthesized using Xilinx® Vivado® for the ZC706 device and met a timing requirement of over 200 MHz. This table shows the resource utilization for the HDL subsystem.

```
% ============================================================
% |Model Name              ||         ImageRotationHDL     ||
% ============================================================
% |Input Image Resolution  ||          480 x 640           ||
% |LUT                     ||             2238             ||
% |FF                      ||             2570             ||
% |BRAM                    ||              96              ||
% |Total DSP Blocks        ||              94              ||
% ============================================================
```

Image Normalization Using External Memory

This example shows how to normalize image pixel values using external memory. The example includes two models that show two ways to model the external memory: SoC external memory modeling and behavioral memory modeling. The example also verifies that the results of the two memory models are the same.

Supported Hardware Platform

- Xilinx® Zynq® ZC706 evaluation kit for the `ImageNormalizationHDLExample` model
- Xilinx® Zynq® ZC706 evaluation kit and FMC-HDMI-CAM mezzanine card for the `soc_imageNormalization_top` model

Introduction

The image normalization algorithm is a preprocessing step in deployment of deep learning networks on FPGA. This example provides an environment to prototype, customize, and integrate an end-to-end application in Simulink®, including a framework for memory-based system integration. The normalization algorithm that is implemented in this example takes reference from the `rescale` function.

In both models, the image normalization algorithm has these inputs and parameters.

- Input image: The image must be in RGB format, with pixels of `uint8` data type.
- Lower bound and upper bound: These values are the range of the normalized output values. These values must be scalars in the range 0 to 255.
- Input minimum and maximum: These values are the minimum and maximum of the input pixel values. You can provide these parameters on the subsystem mask, or you can select the **Compute input minimum and maximum** parameter to automatically calculate these values.

This figure shows the subsystem mask parameters when you clear the **Compute input minimum and maximum** parameter and use fixed values for the **Input minimum** and **Input maximum** parameters.
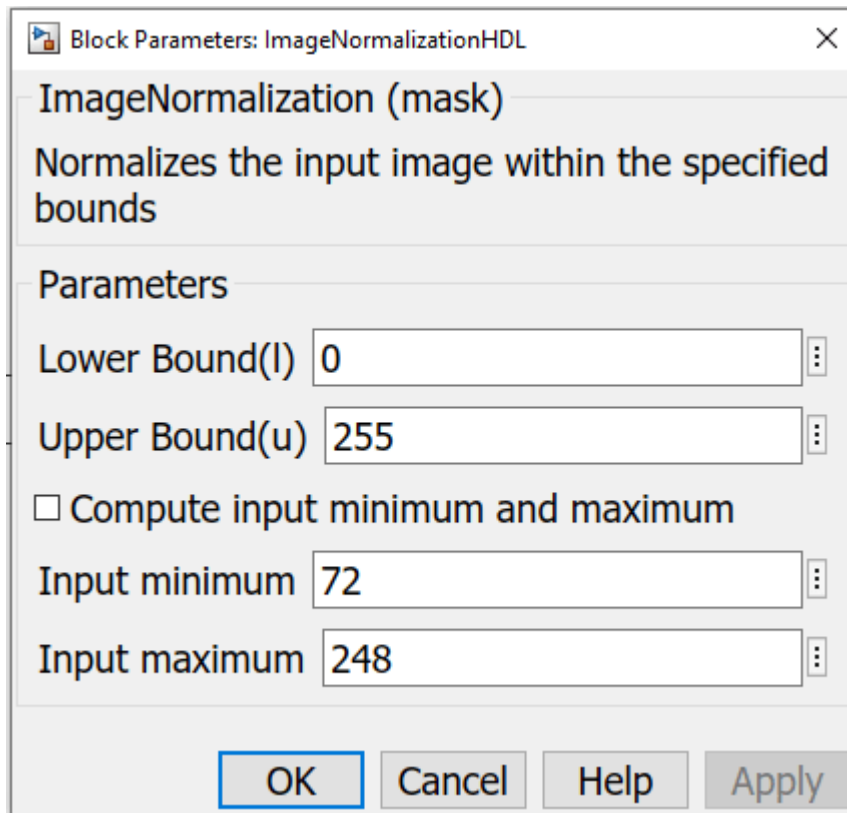
This figure shows the subsystem mask parameters when you select the **Compute input minimum and maximum** parameter. The subsystem computes the input minimum and maximum values from the input pixel stream.
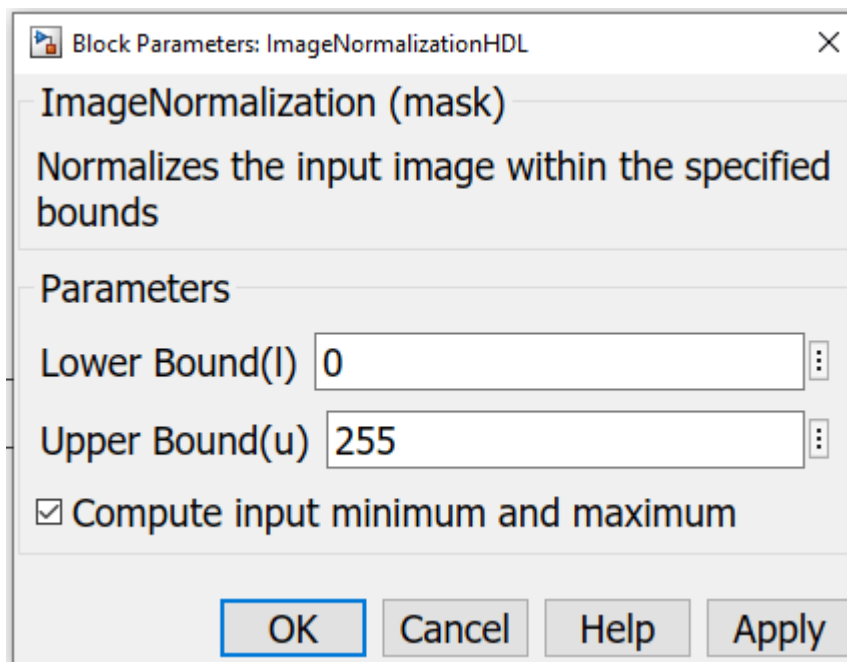
To dynamically calculate the input minimum and maximum of the input frame, the design must store a complete frame in memory. This example shows two ways to model the frame memory. The `ImageNormalizationHDLExample` model stores the input frame by using HDL Coder™ FIFO blocks as a behavioral memory model. The `soc_imageNormalization_top` model stores the input frame by using the SoC Blockset™ AXI4 Random Access Memory block. Using external memory reduces the use of BRAM and enables processing of higher resolution input video streams. The use of external memory requires using AXI4 protocols and verification against memory contention. The model shows a fully compliant AXI4 interface that includes AXI4 write and read controllers.

The AXI4 random access interface provides a simple, direct interface to the memory interconnect. This protocol enables the algorithm to act as a memory controller by providing the addresses and managing the burst transfer directly. The AXI4-Master Write Controller and AXI4-Master Read Controller blocks in this example model a simplified AXI4 interface in Simulink™. When you generate HDL code using the HDL Coder product, the generated code includes a fully compliant AXI4 interface IP.

External Memory Model

The SoC Blockset product provides Simulink blocks and visualization tools for modeling, simulating, and analyzing hardware and software architectures for ASICs, FPGAs, and SoCs. The product enables you to build a system architecture using memory models, bus models, and I/O models, and to simulate the architecture together with the algorithms. This example models external memory using the AXI4 Random Access Memory block from the SoC Blockset library. This block models the connection with hardware through external memory. Both the writer and the reader are managers, sending read and write requests to memory through this block. This block also logs and displays memory performance data. This feature enables you to analyze and debug the performance of the system at simulation time.

HDL Implementation

This figure shows the top level of the `soc_imageNormalization_top` model. The HDMI Rx block processes the video input and passes it to the `soc_imageNormalization_FPGA` reference model.

```
open_system('soc_imageNormalization_top')
```

Image Normalization Using External Memory
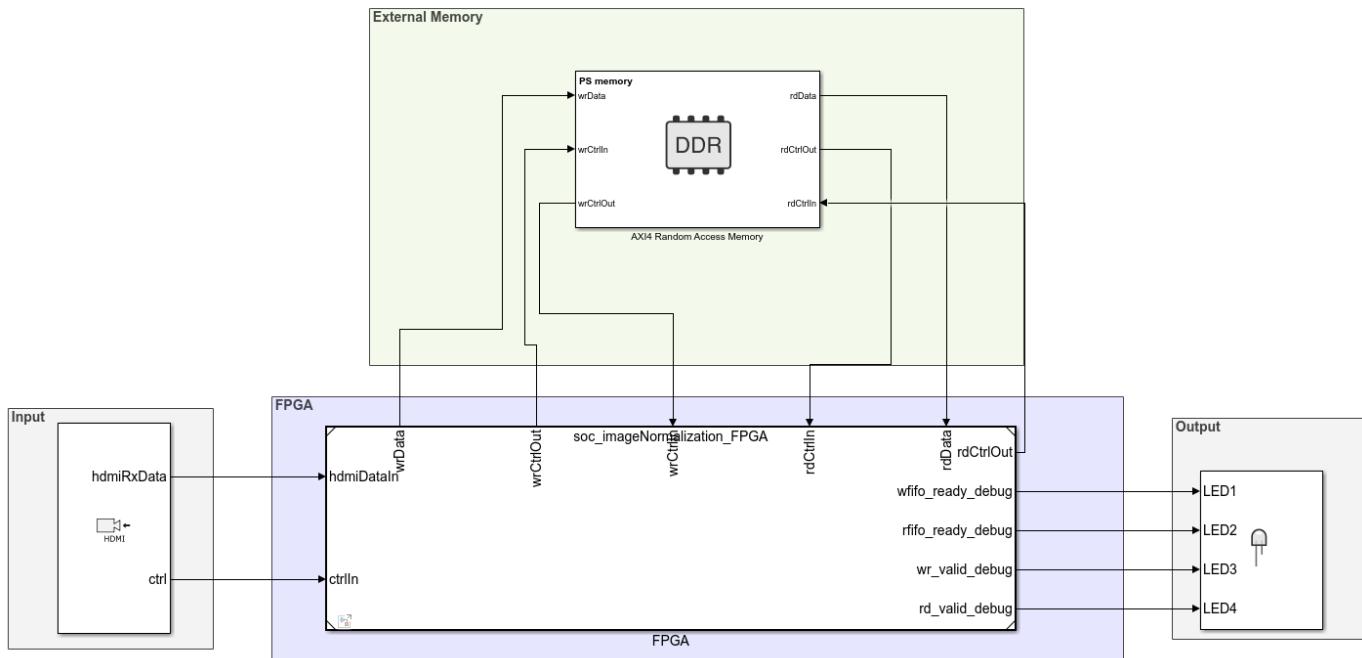
Copyright 2021-2023 The MathWorks, Inc.

In the `soc_imageNormalization_FPGA` model, the input pixel stream connects to a Video Stream Connector block. This block provides a video streaming interface to connect any two IPs in the FPGA implementation. The Video Stream Connector blocks connect the HDMI input and output blocks with the rest of the FPGA algorithm.

```
open_system('soc_imageNormalization_FPGA')
```
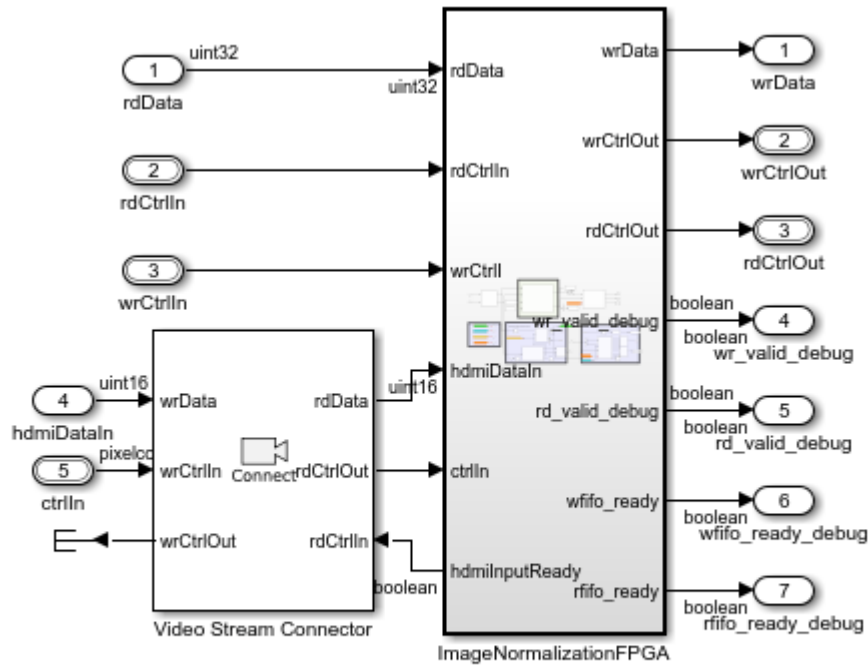
Copyright 2021 The MathWorks, Inc.

The next figure shows the `ImageNormalizationFPGA` subsystem, which implements the AXI write and read from external memory and the normalization algorithm.

The `hdmiDataIn` signal is in YCbCr 4:2:2 pixel stream format. Because the normalization algorithm expects RGB images, the `YCbCr422ToRGB` subsystem converts the YCbCr 4:2:2 data to RGB.

The subsystem contains the `ImageNormalization` subsystem and these sections.

- AXI Write to Memory: This section writes the input data into the memory. It consists of an AXI4-Master Write Controller block that receives the input video control information from the HDMI Rx block and models the AXI4 memory-mapped interface for writing data into the DDR. It has five output signals: `wr_addr`, `wr_len`, `wr_valid`, `rd_start`, and `frame`. The `wr_valid` signal is an input to the AXI Write FIFO block, which stores the incoming pixel intensities. The SoC Bus Creator block generates the `wrCtrlOut` bus for writing the data into the DDR. The model writes one line of data per burst. After writing all of the lines of the frame, the model asserts the `rd_start` signal to begin the read request.

- AXI Read from Memory: This section reads the data from the memory. It consists of an AXI4-Master Read Controller block that receives the `rd_start` signal from the AXI4-Master Write Controller block. The AXI4-Master Read Controller block generates the `rd_addr`, `rd_len`, `rd_avalid`, and `rd_dready` signals. An SoC Bus Creator block combines these signals into a bus. The AXI4-Master Read Controller block also generates the `pixelcontrol` bus corresponding to the `rd_data` signal. The model slices the 32 bit `rd_data` signal to retrieve the 24 bit (LSB) RGB data. Then, the model forms a 1-by-3 `uint8` RGB vector and passes the vector to the normalization algorithm.

The RGB pixel values read from the DDR frame memory are connected to the **buffPixIn** and **buffCtrlIn** input ports of the `Image Normalization` subsystem.

```
open_system('soc_imageNormalization_FPGA/ImageNormalizationFPGA')
```
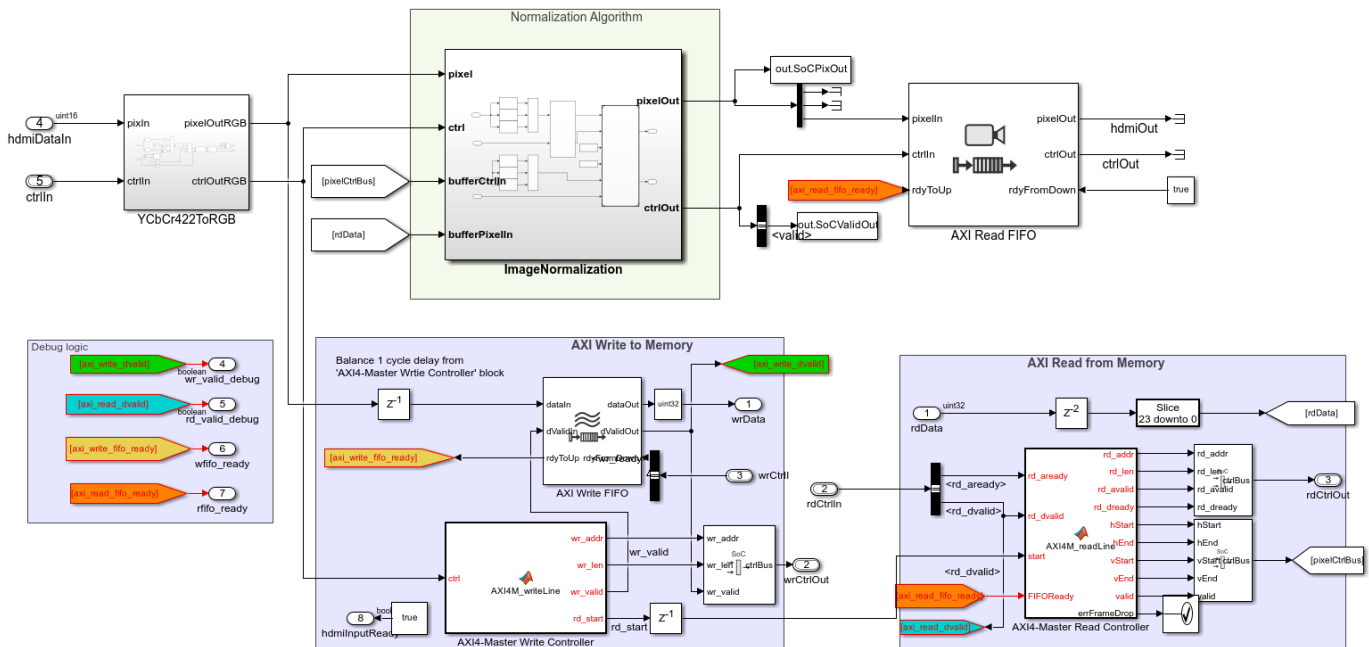

Normalization Algorithm

The next figure shows the `ImageNormalization` subsystem, which implements the normalization algorithm.

The input RGB pixel data (from the `YCbCr422ToRGB` subsystem) is of `ufix24` data type. This subsystem converts the RGB data to `uint8` 1-by-3 RGB vectors. The `InputMinMaxCalc` subsystem calculates the input minimum and maximum values.

The `Rescale` subsystem references the `NormalizationAlgorithm` model.

```
open_system('soc_imageNormalization_FPGA/ImageNormalizationFPGA/ImageNormalization')
```
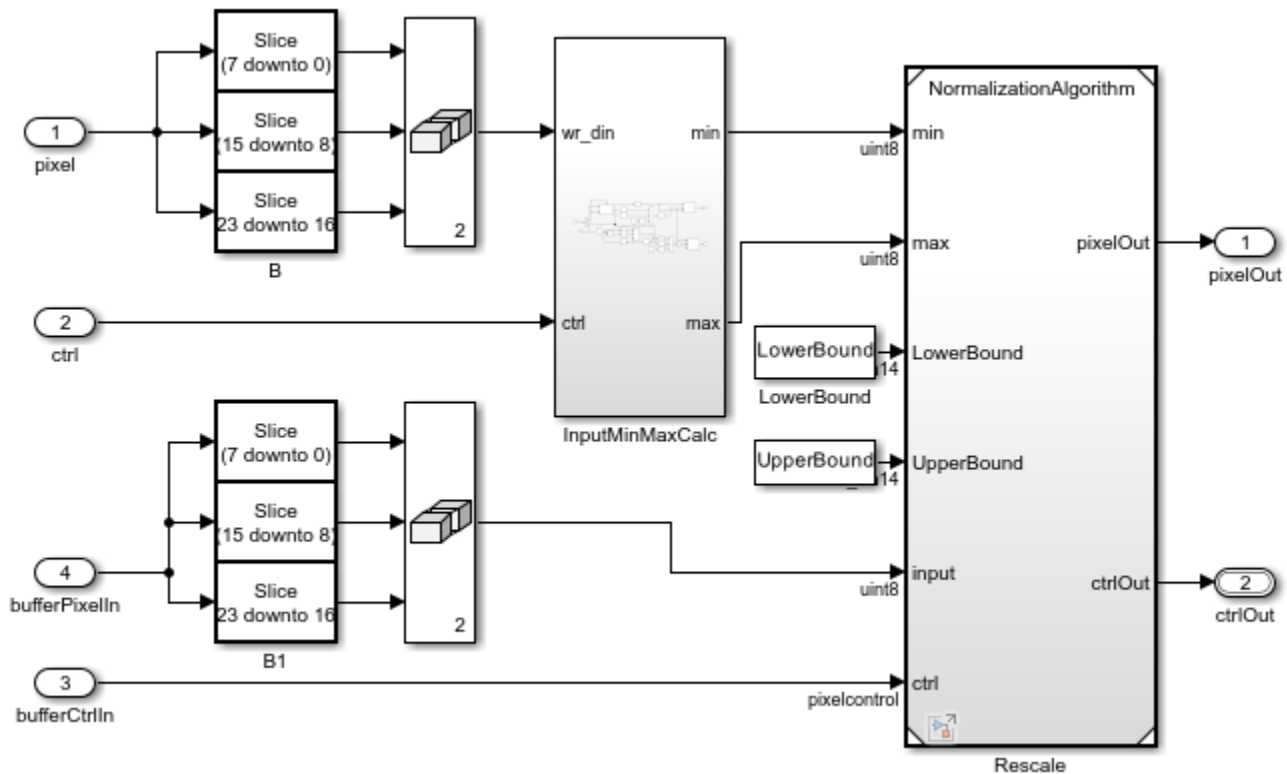
The `NormalizationAlgorithm` model performs the normalization algorithm described by this equation.

$$\text{output} = \frac{(l - u) * (\text{input} - \text{sigma}) + (l * \text{inputMax} - u * \text{inputMin} + l * \text{constReg})}{\text{inputMax} - \text{inputMin} + \text{constReg}}$$

l is the lower bound, u is the upper bound, *sigma* is $\max\left(\min\left(0, \text{inputMax}\right), \text{inputMin}\right)$, and *constReg* is high when the input minimum is equal to the input maximum.

This figure shows the `NormalizationAlgorithm` model.

```
open_system('NormalizationAlgorithm')
```

Copyright 2021 The MathWorks, Inc.
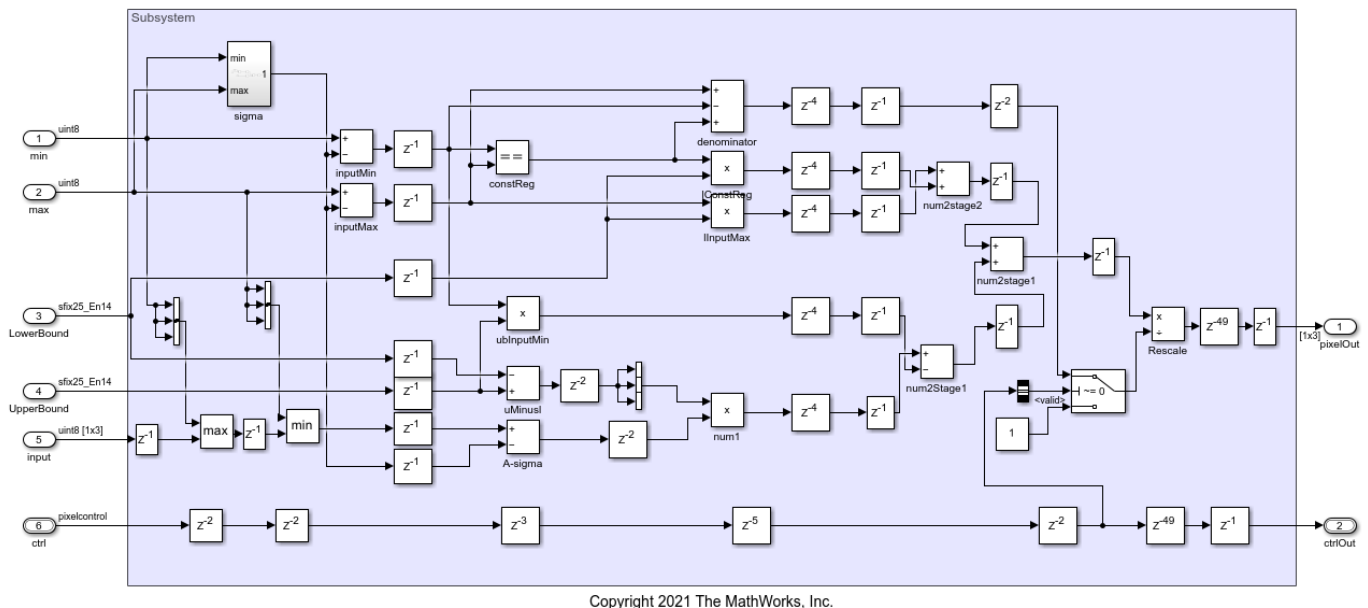
Hardware Implementation

To build, load, and execute the model on FPGA boards, use the **SoC Builder** tool. This example uses the Xilinx Zynq ZC706 evaluation kit. For more detail about the building steps, see SoC Builder (SoC Blockset).

Performance Plots

This example uses an input video of size 480-by-640 pixels. The model configures the HDMI Rx block to use this size. For the Xilinx Zynq ZC706 evaluation kit, the PL DDR controller is configured with a 64 bit AXI4 subordinate interface running at 200 MHz. The resulting bandwidth is 1600 MB/s. This example has two AXI managers connected to the DDR controller. These AXI managers are the AXI4 read and write interfaces of the normalization algorithm. The YCbCr 4:2:2 video format requires 2 bytes per pixel. For the AXI4 read and write interfaces, each pixel is zero-padded to 4 bytes. In this case, the read and write interfaces have a throughput requirement of 2x4x480x640x60 = 147.456 MB/s.

This figure shows the performance plot of the AXI4 Random Access Memory block. To view the performance plot, first open the AXI4 Random Access Memory block. Then, on the **Performance** tab, click **View performance plots**. Select all of the masters under **Bandwidth**, and then click **Update**. After the algorithm starts writing and reading data into external memory, the throughput remains around 180 MB/s, which is within the required throughput of 147.456 MB/s.

Performance Plots for soc_imageNormalization_top/Memory Controller
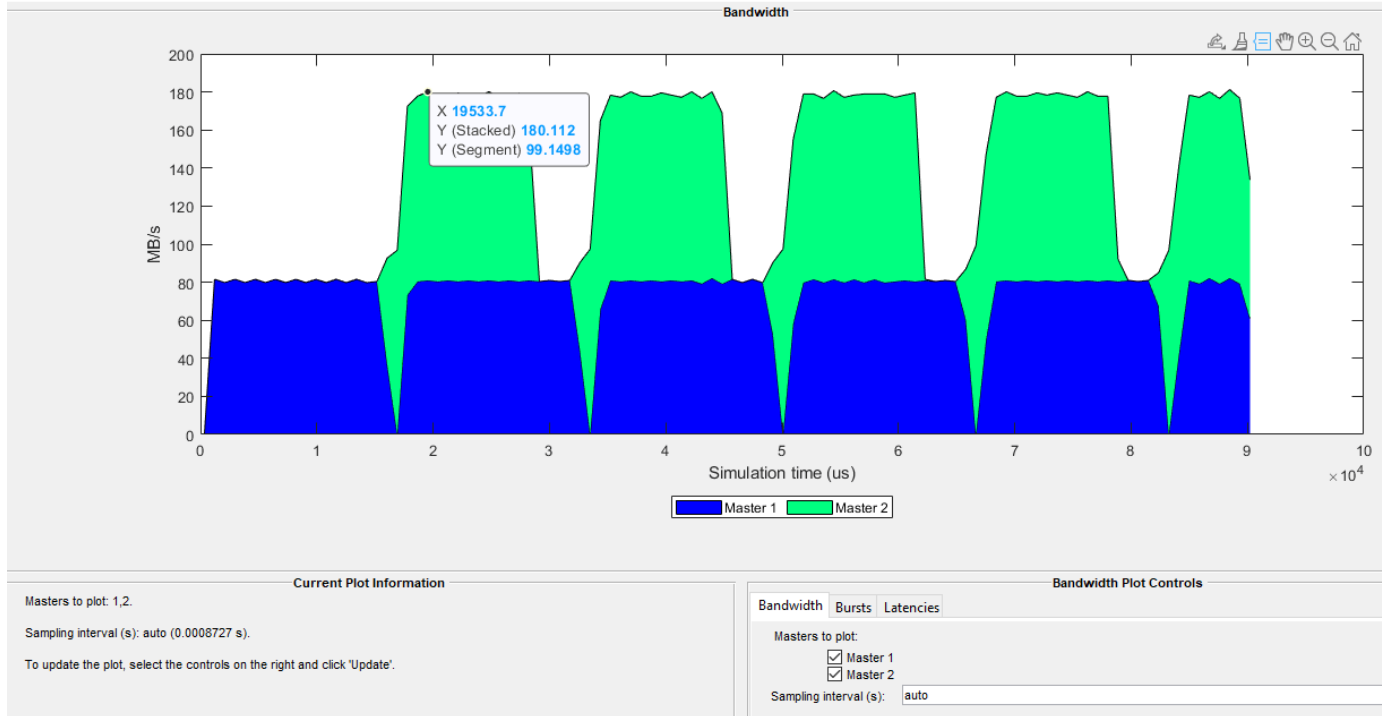
Behavioral Memory Model

This model implements the algorithm using a streaming pixel format, Vision HDL Toolbox™ blocks, and Simulink blocks that support HDL code generation. The serial interface mimics a real-time system and is efficient for hardware designs because less memory is required to store pixel data for computation. The serial interface also enables the design to operate independently of image size and format and makes the design more resilient to timing errors. Fixed-point data types use fewer resources and can give better performance on FPGA. The `InitFcn` callback function initializes the necessary variables for this example.

```
open_system('ImageNormalizationHDLExample');
```
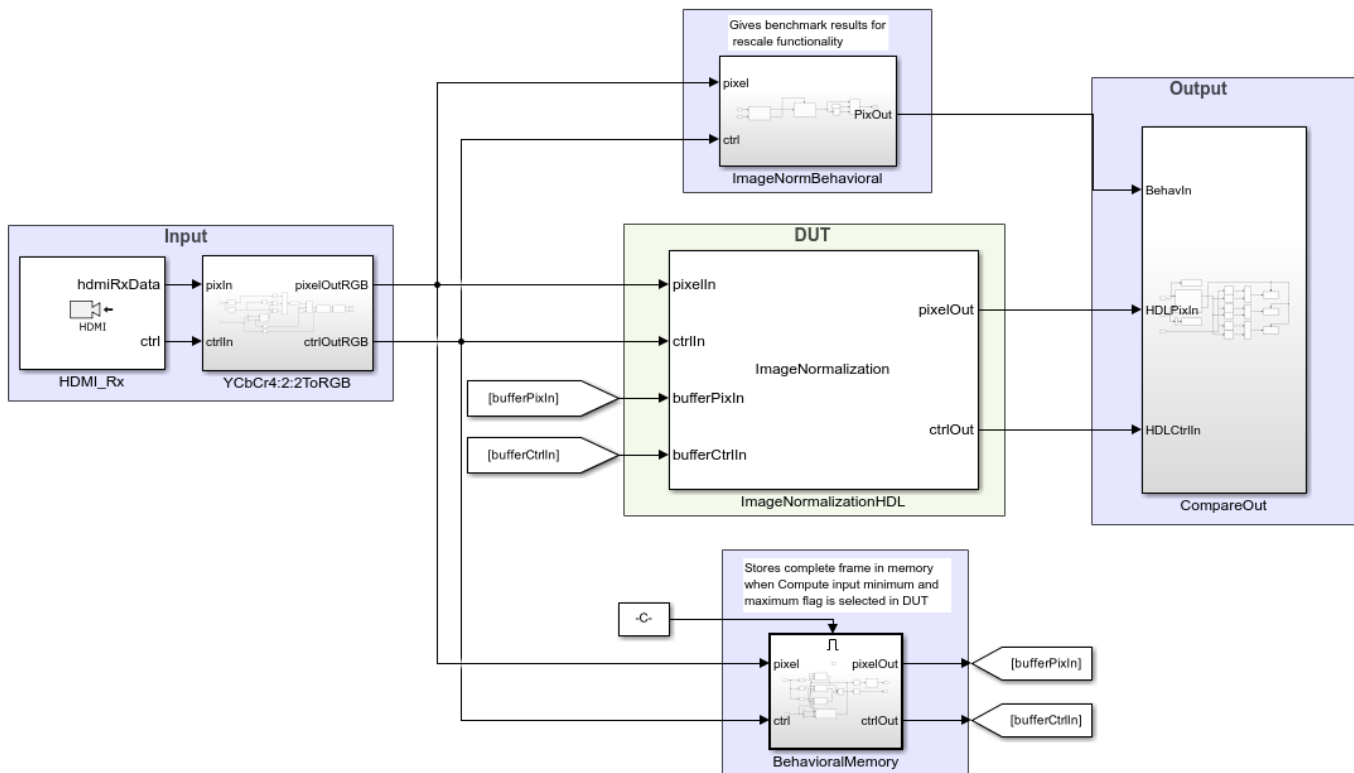
Copyright 2021 The MathWorks, Inc.

The HDMI_Rx block imports the input video to the model. The Pixels To Frame block converts the pixel stream back to image frames. The `BehavioralMemory` subsystem stores the input image so that the `NormalizationAlgorithm` subsystem can read it as needed.

The `ImageNormalizationHDL` subsystem is a variant subsystem that provides either of the two implementations shown in this figure.

```
open_system('ImageNormalizationHDLExample/ImageNormalizationHDL/Variant Subsystem')
```
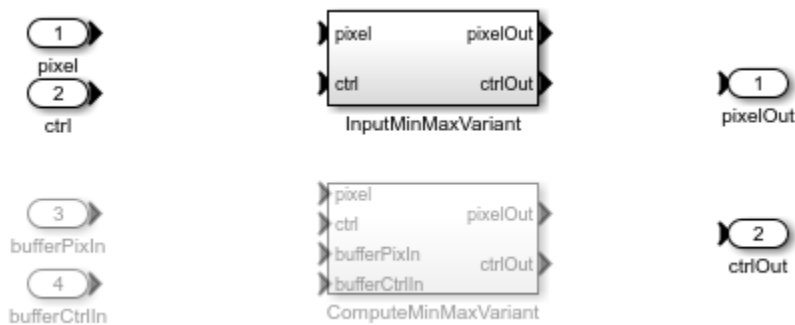

InputMinMaxVariant

If you clear the **Compute input minimum and maximum** parameter, then you must provide **Input minimum** and **Input maximum** parameter values. The algorithm normalizes the input frame by using the provided input minimum and maximum values and the lower and upper bound values.

```
open_system('ImageNormalizationHDLExample/ImageNormalizationHDL/Variant Subsystem/InputMinMaxVari
```
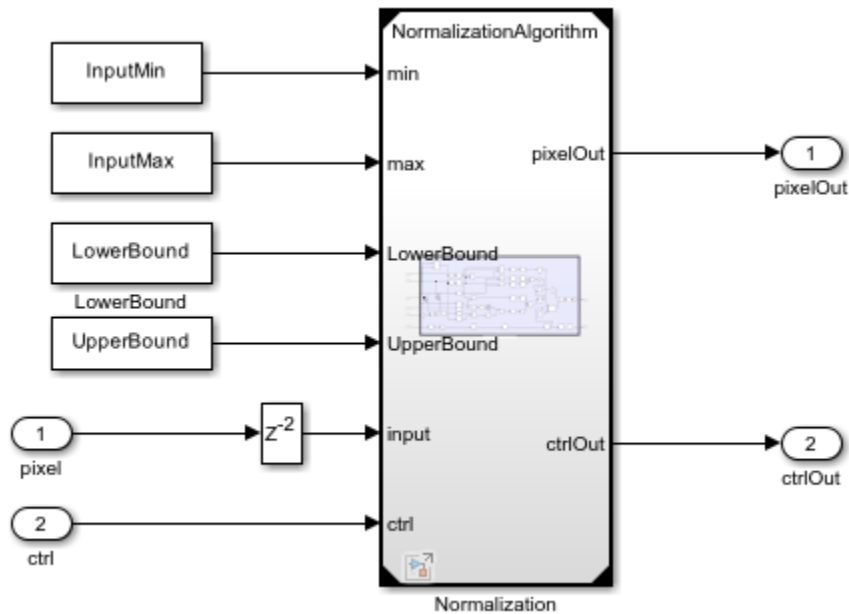

ComputeMinMaxVariant

If you select the **Compute input minimum and maximum** parameter, then the `InputMinMaxCalc` subsystem computes the input minimum and maximum values of the input image. The algorithm normalizes the input frame by using the computed input minimum and maximum values and the provided lower and upper bound values.
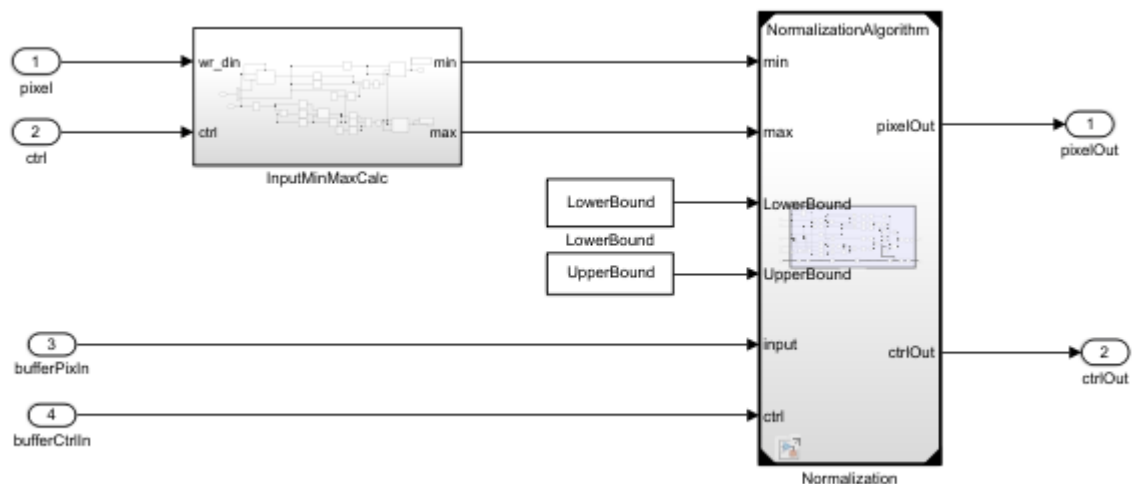
You can verify the results from either of the variant implementations against the golden reference normalization algorithm by using the CompareOut block.

```
open_system('ImageNormalizationHDLExample/CompareOut')
```
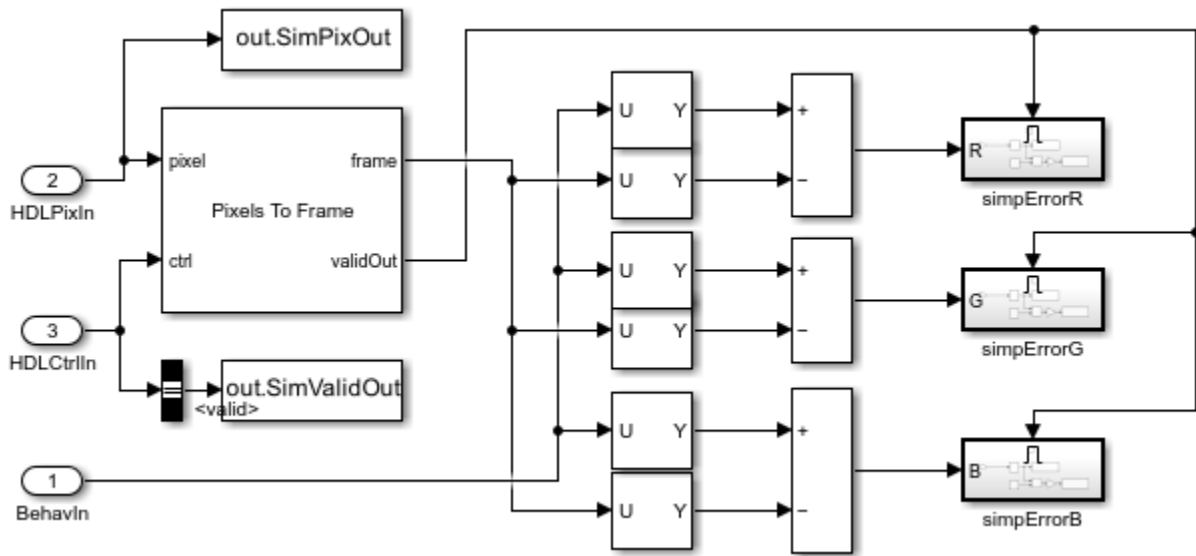
Verify Results Between External Memory Model and Behavioral Memory Model

Compare the output from the `ImageNormalizationHDLExample` model (behavioral memory model) with the output of the `soc_imageNormalization_top` model (external memory model) by using the `errorCheck.m` script. To be able to compare the results of these two models, you must select the **Compute input minimum and maximum** parameter in the `ImageNormalizationHDLExample` model. Run both models to save the output to the MATLAB® workspace. The outputs of the `ImageNormalizationHDLExample` model are the `simPixOut` and `simValidOut` variables. The outputs of the `soc_imageNormalization_top` model are the `socPixOut` and `socValidOut` variables. The `errorCheck` function takes these variables as inputs and returns the total number of error pixels in the R, G, and B channels.

```
[errR,errG,errB] = errorCheck(simPixOut,simValidOut,socPixOut,socValidOut)
```

Contrast Limited Adaptive Histogram Equalization with External Memory

This example shows how to implement the contrast-limited adaptive histogram equalization (CLAHE) algorithm for FPGA, including an external memory interface.

Supported Hardware

• Xilinx® Zynq® ZC706 evaluation kit + FMC-HDMI-CAM mezzanine card

Introduction

Video processing algorithms often store a full frame of video data in memory. Implementing this storage on an FPGA increases BRAM utilization and can result in input video resolution constraints. This example shows how to implement vision algorithms on FPGAs by using an external memory resource to reduce use of BRAM and enable processing of higher resolution input video.

The external memory interface in this example uses AXI4 protocols and verifies the design against memory contention. The AXI4 Random Access interface provides a simple, direct interface to the memory interconnect. This protocol enables the algorithm to act as a memory master by providing the addresses and managing the burst transfer directly. The AXI4 Master Write Controller and AXI4 Master Read Controller blocks in this example model a simplified AXI-4 interface in Simulink®. When you generate HDL code using the HDL Coder™ product, the generated code includes a fully compliant AXI4 interface IP.
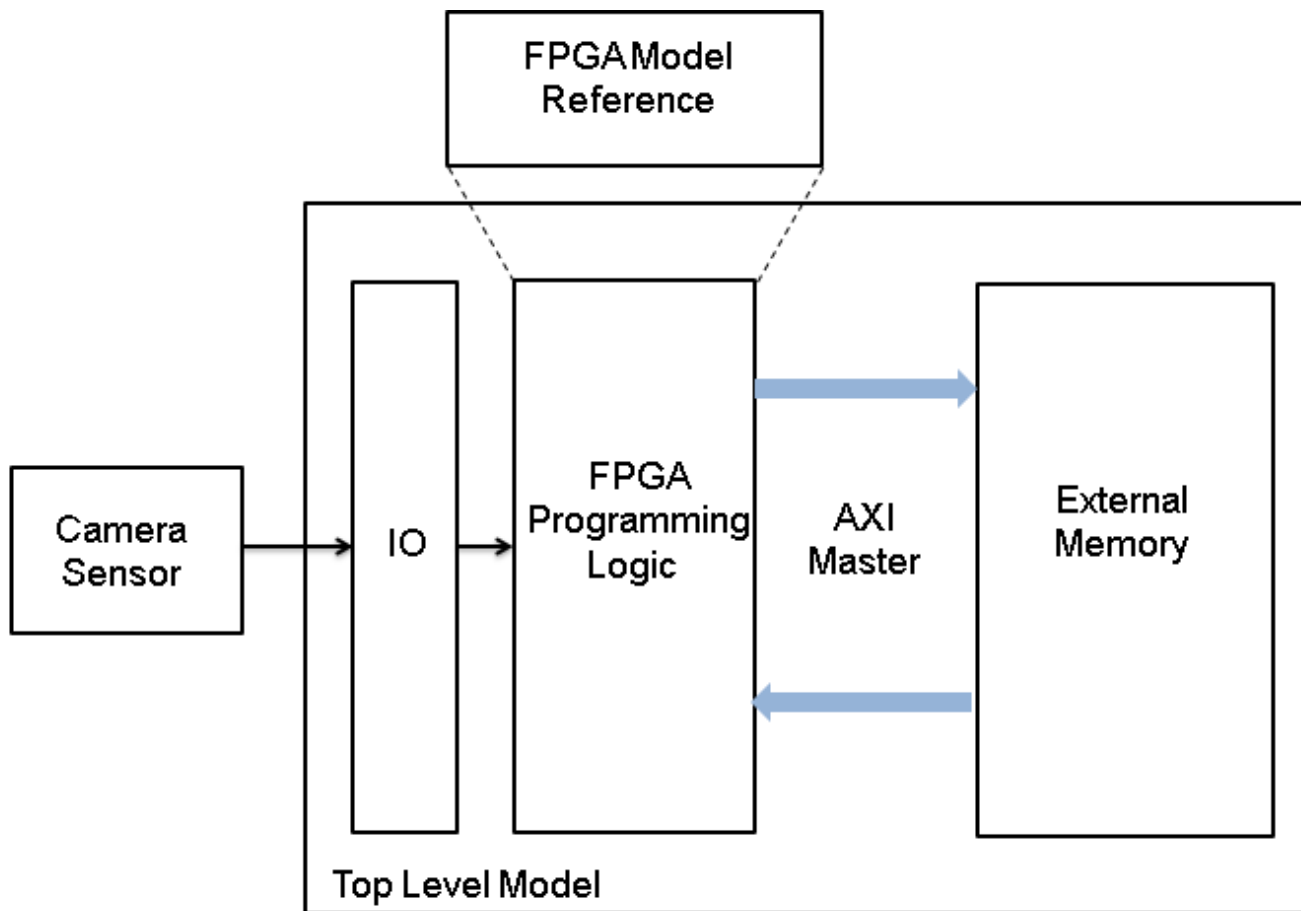
Model External Memory

You can use SoC Blockset™ blocks and visualization tools for modeling, simulating, and analyzing hardware and software architectures for ASICs, FPGAs, and systems on a chip (SoC). These features can help you build system architecture using memory models, bus models, and interface models and help you simulate the architecture together with the algorithms. This example models external memory using the AXI4 Random Access Memory block from the SoC Blockset library. This block models the connection with hardware through external memory. Both the writer and the reader are managers, sending read and write requests to memory through this block. This block also logs and displays memory performance data. This feature enables you to analyze and debug the performance of the system at simulation time.

HDL Implementation

The CLAHE algorithm has three steps: tiling, histogram equalization, and bilinear interpolation. The bilinear interpolation step uses the pixel intensities from the input frame. Storing the full input frame of video data until the bilinear interpolation step requires external memory.

The figure shows the top level of the example model. The HDMI Rx block processes the video input and passes it to the CLAHEAlgorithm_fpga subsystem. The HDMI Rx block converts raw video data to a YCbCr 4:2:2 pixel stream format. The output data is a pixel stream suitable for hardware algorithm design. The HDMI Rx block also directs the **SoC Builder** tool to generate the IP blocks necessary to receive video data from the FMC-HDMI-CAM card that is attached to the hardware board.

In the model, the AXI4-Master Write Controller and AXI4-Master Read Controller blocks model the AXI4 memory mapped interfaces. The AXI4-Master Write Controller block writes the input frame into the external memory, and the AXI4-Master Read Controller block reads the frame from the external memory for bilinear interpolation. The AXI Read FIFO block sends the output pixel stream to the HDMI Tx block. The HDMI Tx block converts a pixel stream in YCbCr 4:2:2 format to raw video data for display during simulation. This block also directs the **SoC Builder** tool to generate the IP blocks that transmit video data back to the FMC-HDMI-CAM card. To indicate the status of the AXI Read FIFO and AXI Write FIFO blocks when running the design on hardware, four debug signals from these blocks are connected to LEDs on the board.

CLAHE With External Memory
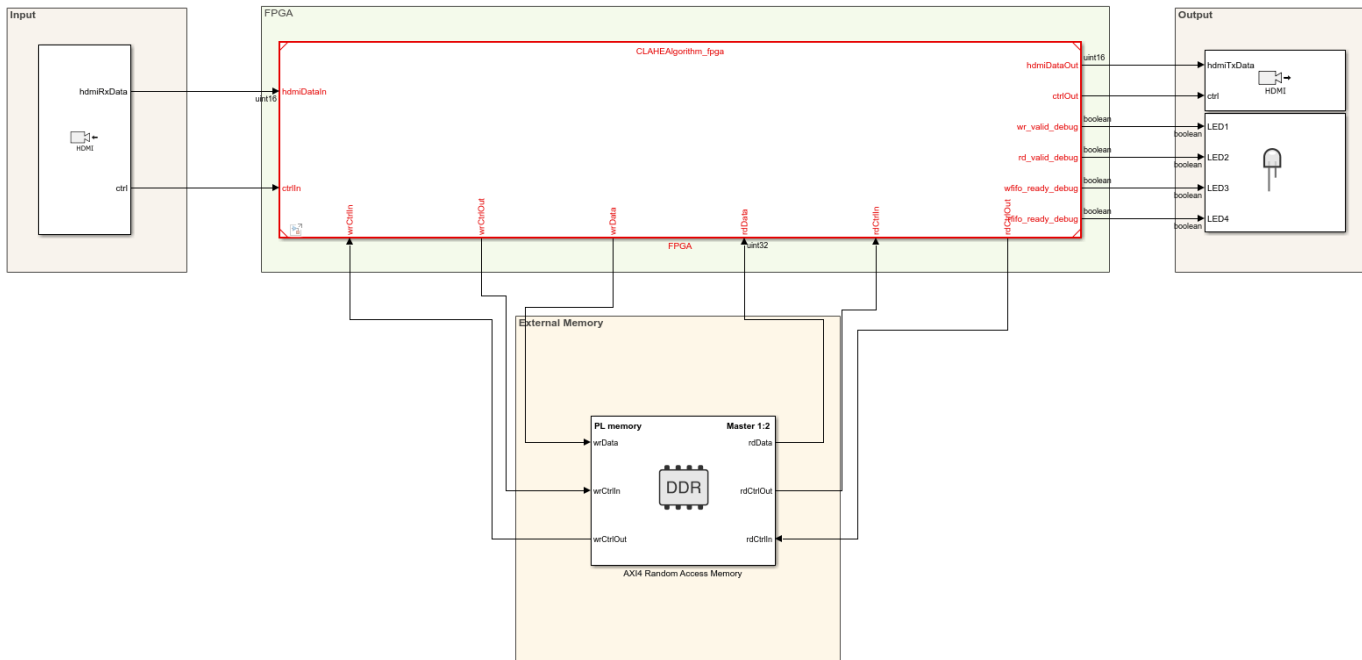
The next figure shows the CLAHEAlgorithm_fpga reference model. The input pixel stream connects to a Video Stream Connector block. This block provides a video streaming interface to connect any two IPs in the FPGA implementation. In this example, the Video Stream Connector blocks connect the HDMI input and output blocks with the rest of the FPGA algorithm.
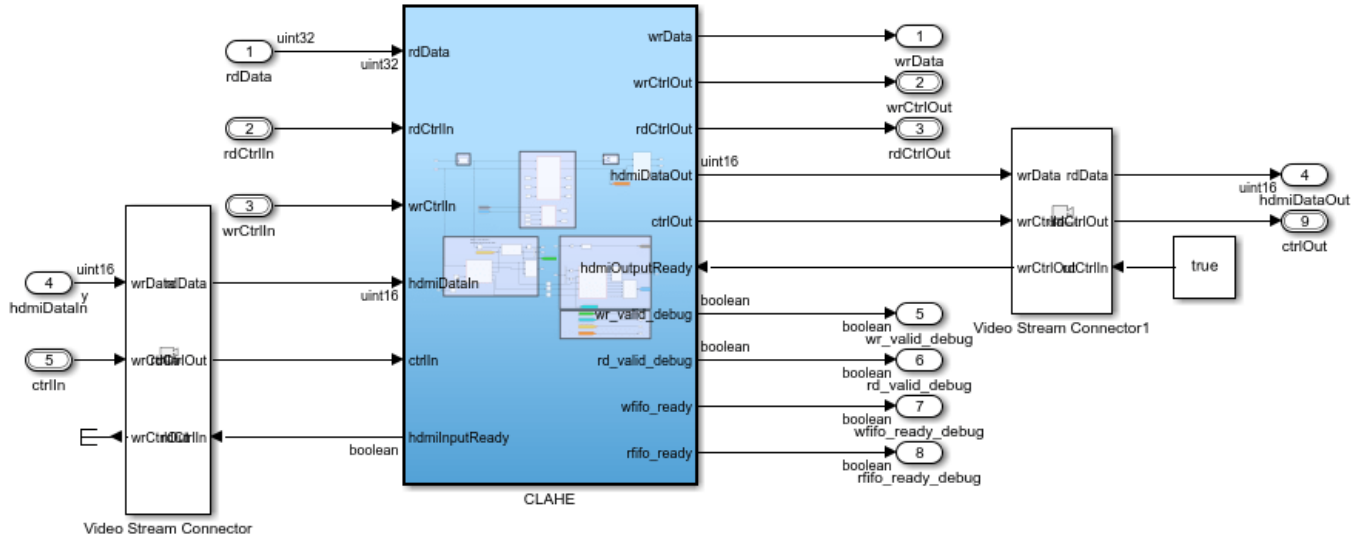
Copyright 2021-2022 The MathWorks, Inc.

The next figure shows the CLAHEAlgorithm_fpga/CLAHE subsystem, which implements the AXI write and read from external memory, and the CLAHE algorithm.
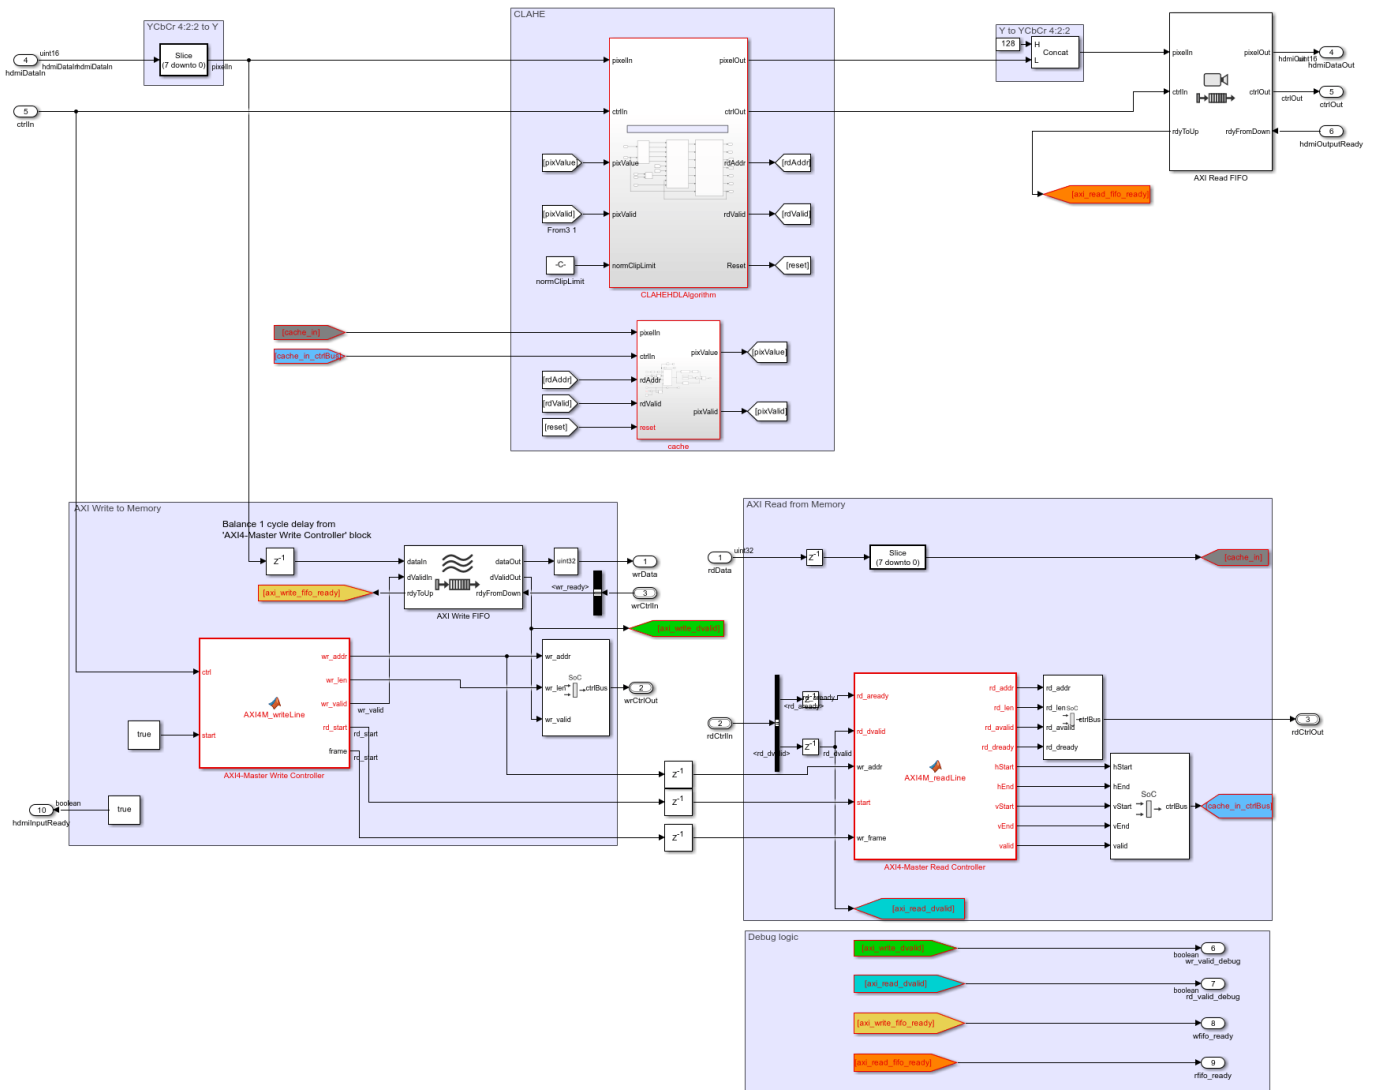
The subsystem contains these areas: * AXI Write to Memory: This section writes the input data into the DDR. It consists of an AXI4 Master Write Controller block that receives the input video control information from the HDMI Rx block and models the AXI4 memory mapped interface for writing data into the DDR. It generates five signals: `wr_addr`, `wr_len`, `wr_valid`, `rd_start`, and `frame`. The `wr_valid` signal is an input to the AXI Write FIFO block, which stores the incoming pixel intensities. The SoC Bus Creator block generates the `wrCtrlOut` master to slave bus for writing the data into the DDR. The model writes one line of data per burst. After writing *tileHeight*/ 2 lines (where *tileHeight* corresponds to the height of each tile in CLAHE), the model asserts the `rd_start` signal to begin the read request. The `frame` signal indicates the input frame count.

- AXI Read from Memory: This section reads the data from the DDR. It consists of an AXI4-Master Read Controller block that receives the `rd_start` signal from the AXI4-Master Write Controller block. The AXI4-Master Read Controller block generates the `rd_addr`, `rd_len`, `rd_avalid`, and `rd_dready` signals. An SoC Bus Creator block combines these signals into a bus. The AXI4-Master Read Controller block also generates the `pixelcontrol` bus corresponding to the `rd_data`. The model slices the 32-bit `rd_data` signal to retrieve the 8-bit (LSB) luminance component and then writes it into the cache memory block of the CLAHE algorithm.

- CLAHE: For a detailed description of the implementation of the CLAHE algorithm for hardware, see the "Contrast Limited Adaptive Histogram Equalization" on page 2-163 example. In this example, the CLAHEHDLAlgorithm subsystem operates on 8-bit grayscale images, which is why the 8-bit luminance (Y) component is separated from the 16-bit YCbCr pixel data.

The CLAHEHDLAlgorithm subsystem performs the three steps of CLAHE: tiling, histogram equalization, and bilinear interpolation. In the first step, the input frame is divided into a grid of tiles. In the second step, the histogram of each tile is calculated, and then performs distribution, redistribution, and CDF calculations. The calculated CDF values are stored in a buffer for further processing. The third step calculates the output pixel intensities by using a bilinear interpolation of the CDF values. The pixel intensities of the input frame are used as the address to the buffer that stores the CDF values. These pixel intensities are read from the external memory that stores the original input frame.

Because the data read back from the external memory is in burst mode, it cannot be used directly for bilinear interpolation. The cache buffer stores the burst of lines read from the external memory. The depth of the cache is enough to store a number of lines equal to *tileHeight*. The rdValid signal from the CLAHEHDLAlgorithm subsystem generates the rd_addr signal to read the data from the cache. The data read from the cache (pixValue) is then returned to the CLAHEHDLAlgorithm subsystem to complete the bilinear interpolation to calculate the output pixel intensity.

Hardware Implementation

The **SoC Builder** tool builds, loads, and executes the model on the FPGA board. The hardware board used in this example is the Xilinx Zynq ZC706 evaluation kit. To build, load, and execute the design on the hardware, follow these steps.

1 Set up the Vivado® tool for synthesis, implementation, and generation of the FPGA bitstream.
2 The example model runs in Accelerator mode by default to speed up the simulation. However, the **SoC Builder** tool requires Normal simulation mode. In Simulink Configuration Parameters, set **Simulation mode** to Normal.
3 Launch the **SoC Builder** tool by clicking **Configure, Build, & Deploy** in the Simulink toolstrip.
4 On the **Setup** screen, select **Build model**. Click **Next**.
5 On the **Select Build Action** screen, select **Build, load, and run**. Click **Next**.
6 On the **Select Project Folder** screen, specify the project folder. Click **Next**.
7 On the **Review Memory Map** screen, to view the memory map, click **View/Edit**. Click **Next**.
8 On the **Validate Model** screen, to check the compatibility of the model for implementation, click **Validate**. Click **Next**.
9 On the **Build Model** screen, to build the model, click **Build**. An external shell opens when FPGA synthesis begins. Click **Next**.
10 When the bitstream generation is complete, on the **Connect Hardware** screen, to test the connectivity between the host computer and the hardware board, click **Test Connection**. Load the bitstream on the hardware by clicking **Load**.

This figure shows the final **SoC Builder** results after these steps are complete.

Simulation and Results

This example uses an input video of size 480-by-640 pixels. This size is configured in the HDMI Rx block. For the Xilinx Zynq ZC706 evaluation kit, the PL DDR controller is configured with a 64-bit AXI4-Slave interface running at 200 MHz. The resulting bandwidth is 1600 MB/s. This example has two AXI masters connected to the DDR controller. These AXI masters are the DUT AXI4 read and write interfaces. The YCbCr 4:2:2 video format requires 2 bytes per pixel. For the DUT AXI4 read and write interfaces, each pixel is zero-padded to 4 bytes. In this case, the read and write interfaces have a throughput requirement of 2*4*480*640*60 = 147.456 MB/s.

This figure shows the performance plot of the AXI4 Random Access Memory block. To view the performance plot, first open the AXI4 Random Access Memory block. Then, on the **Performance** tab, click **View performance plots**. Select all masters under **Bandwidth**, and then click **Update**. After the DUT starts writing and reading data into external memory, the throughput remains around 154 MB/s, which is within the required throughput of 147.456 MB/s.
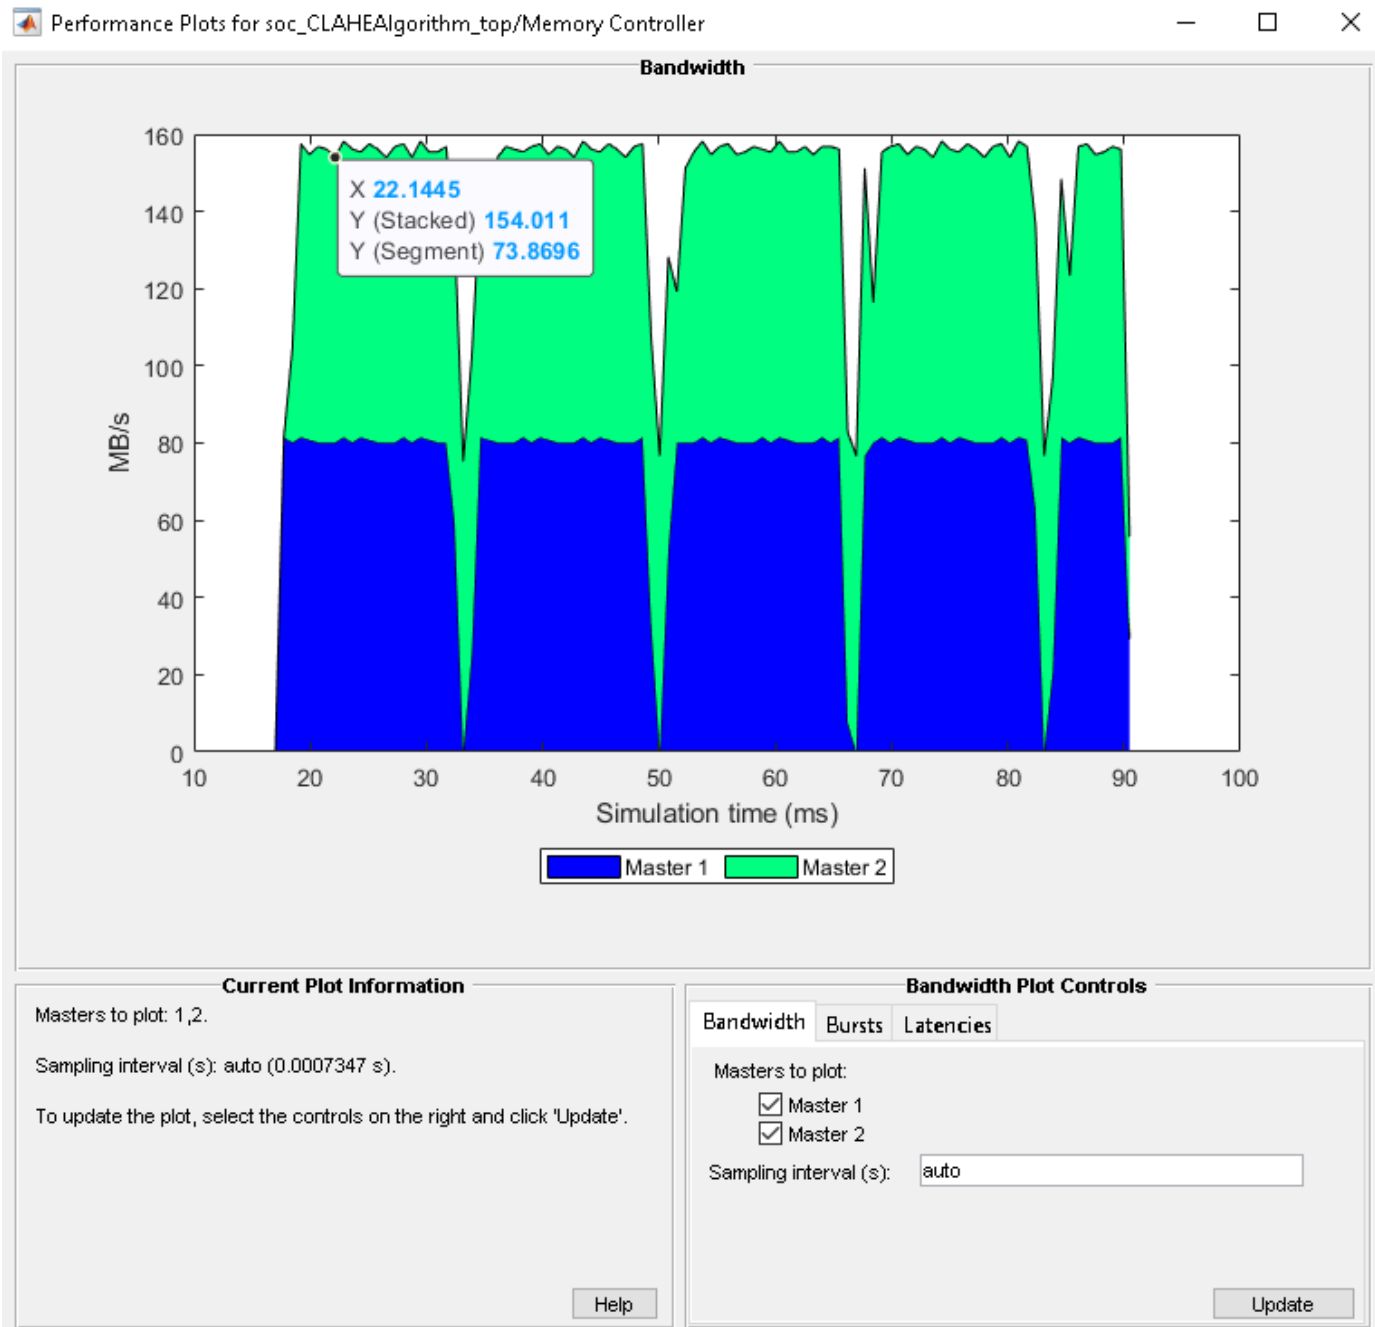
The signals in the example model are logged during simulation. View these signals by using the **Logic Analyzer** app. This figure shows the logged data of input and output frames.
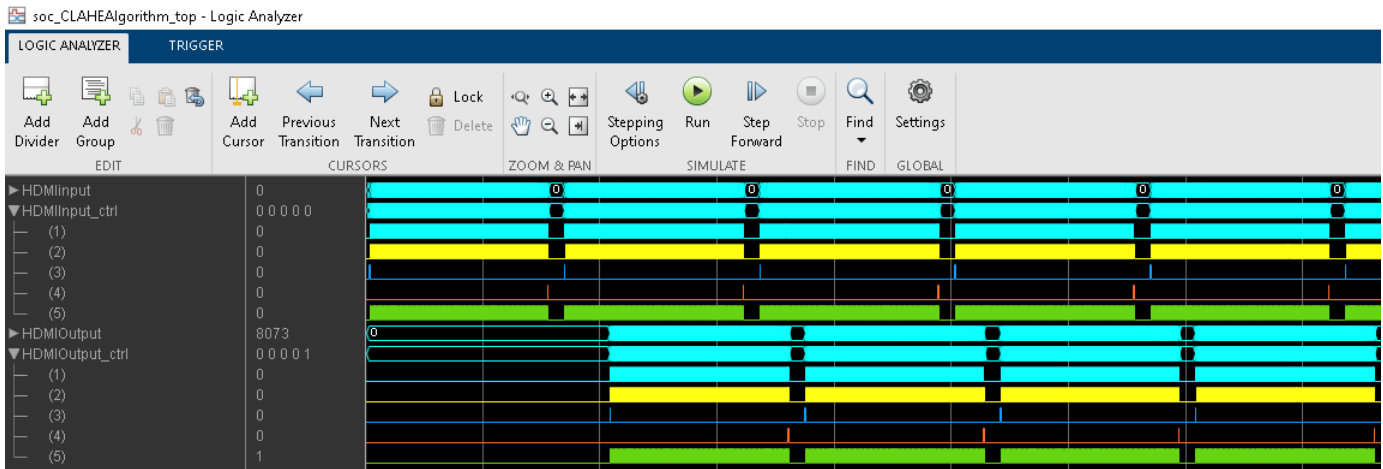
This figure shows the input and output frames from the model. The result shows the improved contrast in the output image.
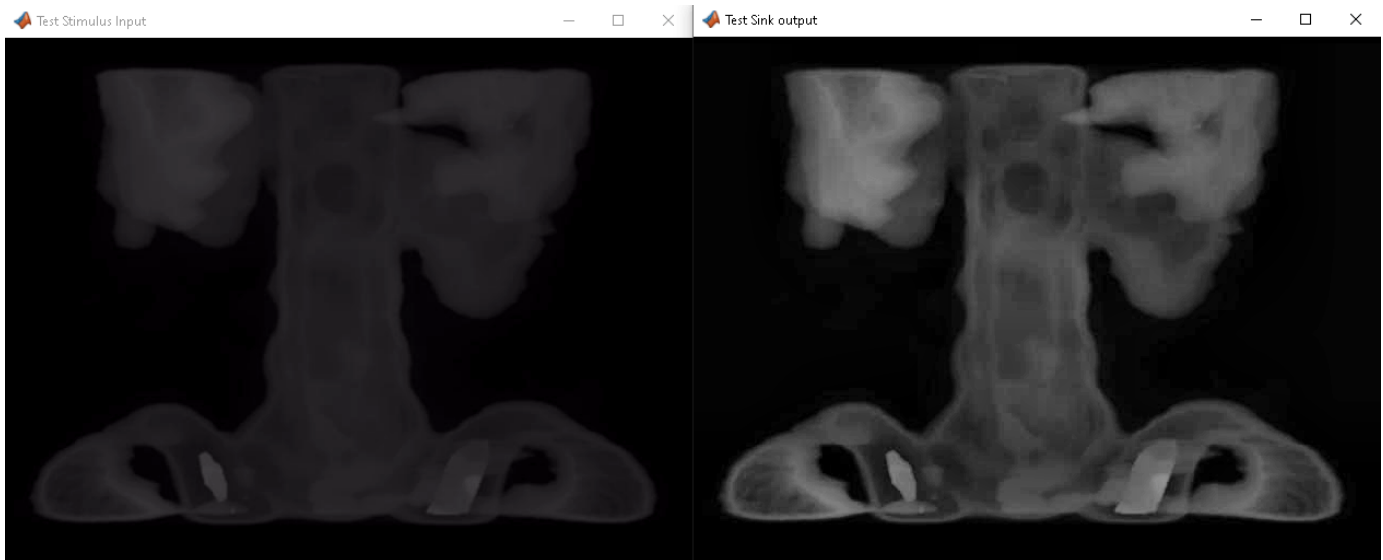
References

[1] Zuiderveld, Karel. "Contrast Limited Adaptive Histogram Equalization." In Graphics Gems IV, edited by Paul S. Heckbert, 474-485. AP Professional, 1994.

See Also

Memory Channel | Memory Controller | Memory Traffic Generator

Related Examples

- "Vertical Video Flipping Using External Memory" on page 3-60

HDL Cosimulation

HDL cosimulation links an HDL simulator with MATLAB or Simulink. This communication link enables integrated verification of the HDL implementation against the design. To perform this integration, you need an HDL Verifier™ license. HDL Verifier cosimulation tools enable you to:

- Use MATLAB or Simulink to create test signals and software test benches for HDL code
- Use MATLAB or Simulink to provide a behavioral model for an HDL simulation
- Use MATLAB analysis and visualization capabilities for real-time insight into an HDL implementation
- Use Simulink to translate legacy HDL descriptions into system-level views

See Also

More About

- "HDL Cosimulation" (HDL Verifier)

FPGA-in-the-Loop

FPGA-in-the-loop (FIL) enables you to run a Simulink or MATLAB simulation that is synchronized with an HDL design running on an FPGA board. This link between the simulator and the board enables you to verify HDL implementations directly against Simulink or MATLAB algorithms. You can apply real-world data and test scenarios from these algorithms to the HDL design that is running on the FPGA.

In Simulink, you can use the FIL Frame To Pixels and FIL Pixels To Frame blocks to accelerate communication between Simulink and the FPGA board. In MATLAB, you can modify the generated code to speed up communication with the FPGA board.

FPGA-in-the-Loop Simulation with Vision HDL Toolbox Blocks

This example shows how to modify the generated FPGA-in-the-loop (FIL) model for more efficient simulation of the Vision HDL Toolbox™ streaming video protocol.

Autogenerated FIL Model

When you generate a programming file for a FIL target in Simulink, the HDL Workflow Advisor creates a model to compare the FIL simulation with your Simulink design. For details of how to generate FIL artifacts for a Simulink model, see "FIL Simulation with HDL Workflow Advisor for Simulink" (HDL Verifier).

For Vision HDL Toolbox designs, the FIL block in the generated model replicates the pixel-streaming interface and sends one pixel at a time to the FPGA. The model shown was generated from the example model in "Design Video Processing Algorithms for HDL in Simulink".
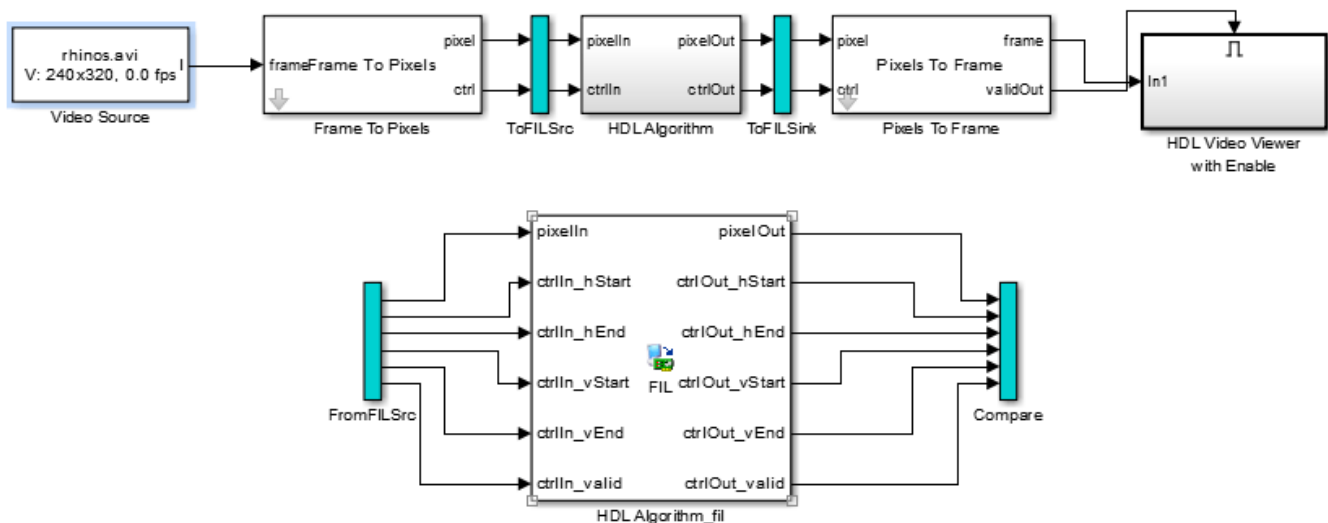
The top part of the model replicates your Simulink design. The generated FIL block at the bottom communicates with the FPGA. ToFILSrc subsystem copies the pixel-stream input of the HDL Algorithm block to the FromFILSrc subsystem. The ToFILSink subsystem copies the pixel-stream output of the HDL Algorithm block into the Compare subsystem, where it is compared with the output of the HDL Algorithm_fil block. For image and video processing, this setup is slow because the model

sends only a single pixel, and its associated control signals, in each packet to and from the FPGA board.

Modified FIL Model for Pixel Streaming

To improve the communication bandwidth with the FPGA board, you can use the generated FIL block with vector input rather than streaming. This example includes a model, FILSimulinkWithVHTExample.slx, created by modifying the generated FIL model. The modified model uses the FIL Frame To Pixels and FIL Pixels To Frame blocks to send one frame at a time to the generated FIL block. You cannot run this model as is. You must generate your own FIL block and bitstream file that use your board and connection settings.

Copyright 2017 The MathWorks, Inc.
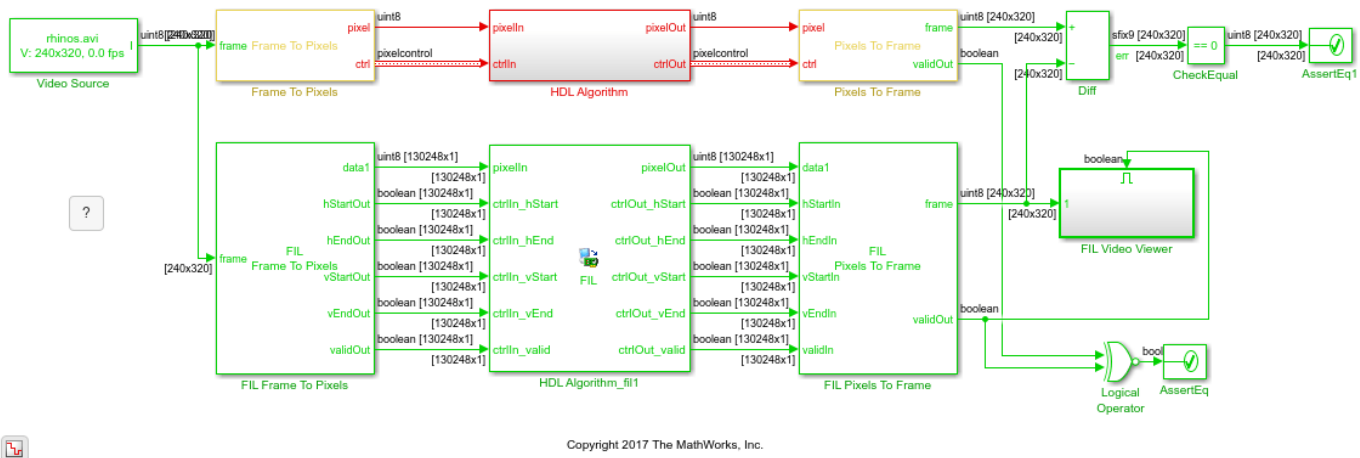
To convert from the generated model to the modified model:

1. Remove the ToFILSrc, FromFILSrc, ToFILSink, and Compare subsystems, and create a branch at the frame input of the Frame To Pixels block.

2. Insert the FIL Frame To Pixels block before the HDL Algorithm_fil block. Insert the FIL Pixels To Frame block after the HDL Algorithm_fil block.

3. Branch the frame output of the Pixels To Frame block for comparison. You can compare the entire frame at once with a Diff block. Compare the `validOut` signals using an XOR block.

4. In the FIL Frame To Pixels and FIL Pixels To Frame blocks, set the Video format parameter to match the video format of the Frame To Pixels and Pixels To Frame blocks.

5. Set the **Vector size** in the FIL Frame To Pixels and FIL Pixels To Frame blocks to `Frame` or `Line`. The size of the FIL Frame To Pixels vector output must match the size of the FIL Pixels To Frame vector input. The vector size of the FIL block interfaces does not modify the generated HDL code. It affects only the packet size of the communication between the simulator and the FPGA board.

The modified model sends an entire frame to the FPGA board in each packet, significantly improving the efficiency of the communication link.

FPGA-in-the-Loop Simulation with Multipixel Streaming

When using FPGA-in-the-Loop with a multipixel streaming design, you must flatten the pixel ports to vectors for input and output of the FIL block. Use Selector blocks to separate the input pixel streams into *NumPixels* vectors, and use a Vector Concatenate block to recombine the output vectors.

If each pixel is represented by more than one component, the FIL Frame To Pixels block has one data port per component and the FIL block has *NumPixels×NumComponents* ports. Split each component matrix into *NumPixels* vectors.

This model shows a multipixel, single component design.
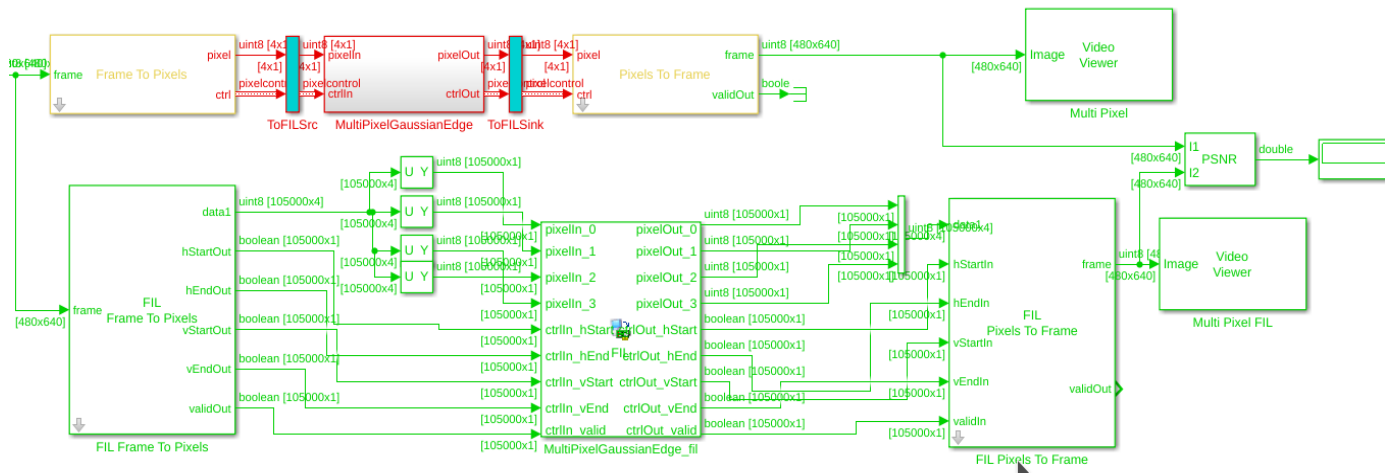
For VHDL code generation, in **Configuration Parameters > HDL Code Generation > Global Settings > Ports**, set the **Scalarize ports** parameter to DUT Level.
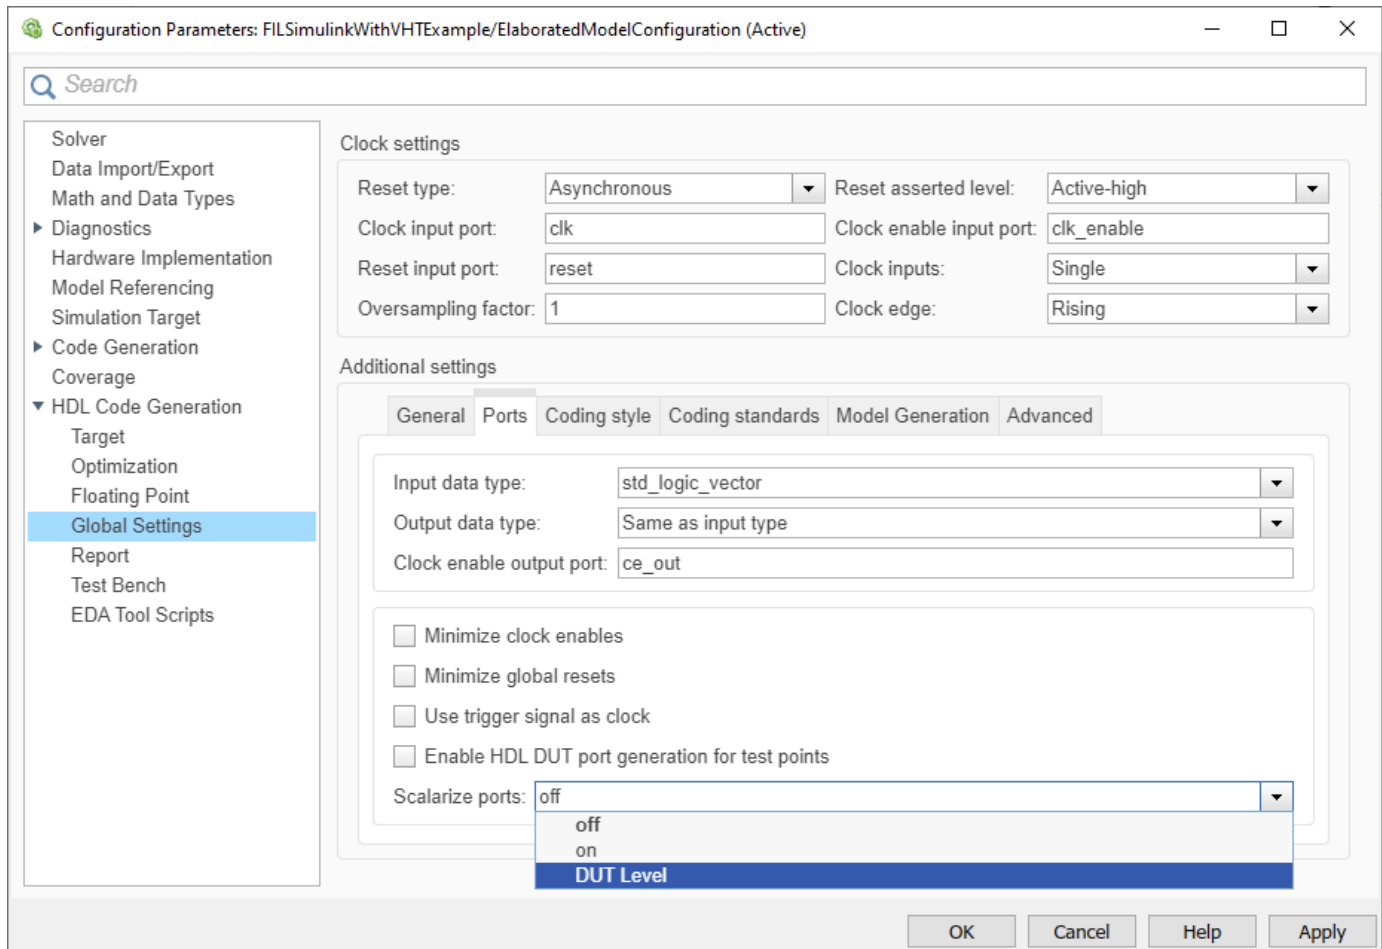
FPGA-in-the-Loop Simulation with Vision HDL Toolbox System Objects

This example shows how to modify the generated FPGA-in-the-loop (FIL) script for more efficient simulation of the Vision HDL Toolbox™ streaming video protocol. For details of how to generate FIL artifacts for a MATLAB® System object™, see "FIL Simulation with HDL Workflow Advisor for MATLAB" (HDL Verifier).

Autogenerated FIL Function

When you generate a programming file for a FIL target in MATLAB, the HDL Workflow Advisor creates a test bench to compare the FIL simulation with your MATLAB design. For Vision HDL Toolbox designs, the *DUTname*_fil function in the test bench replicates the pixel-streaming interface and sends one pixel at a time to the FPGA. *DUTname* is the name of the function that you generated HDL code from.

This code snippet is from the generated test bench *TBname*_fil.m, generated from the example script in "Pixel-Streaming Design in MATLAB" on page 2-206. The code calls the generated *DUTname*_fil function once for each pixel in a frame.

```
for p = 1:numPixPerFrm
    [pixOutVec( p ),ctrlOutVec( p )] = PixelStreamingDesignHDLDesign_fil( pixInVec( p ), ctrlInVe
end
```

The generated *DUTname*_fil function calls your HDL-targeted function. It also calls the
*DUTname*_sysobj_fil function, which contains a System object that connects to the FPGA.
*DUTname*_fil compares the output of the two functions to verify that the FPGA implementation
matches the original MATLAB results. This snippet is from the file *DUTname*_fil.m.

```
% Call the original MATLAB function to get reference signal
[ref_pixOut,tmp_ctrlOut] = PixelStreamingDesignHDLDesign(pixIn,ctrlIn);

   ...

% Run FPGA-in-the-Loop
[pixOut,ctrlOut_hStart,ctrlOut_hEnd,ctrlOut_vStart,ctrlOut_vEnd,ctrlOut_valid] ...
  = PixelStreamingDesignHDLDesign_sysobj_fil(pixIn,ctrlIn_hStart,ctrlIn_hEnd,ctrlIn_vStart,ctrlIn

   ...

% Verify the FPGA-in-the-Loop output
hdlverifier.assert(pixOut,ref_pixOut,'pixOut');
```

For image and video processing, this setup is slow because the function sends only one pixel, and its
associated control signals, in each packet to and from the FPGA board.

Modified FIL Test Bench for Pixel Streaming

To improve the communication bandwidth with the FPGA board, you can modify the autogenerated
test bench, *TBname*_fil.m. The modified test bench calls the FIL System object directly, with one
frame at a time. These snippets are from the
PixelStreamingDesignHDLTestBench_fil_frame.m script, modified from FIL artifacts
generated from the example script in "Pixel-Streaming Design in MATLAB" on page 2-206. You
cannot run this script as is. You must generate your own FIL System object, function, and bitstream
file that use your board and connection settings. Then, either modify your version of the generated
test bench, or modify this script to use your generated FIL object.

Declare an instance of the generated FIL System object.

```
fil = class_PixelStreamingDesignHDLDesign_sysobj;
```

Comment out the loop over the pixels in the frame.

```
%          for p = 1:numPixPerFrm
%              [pixOutVec( p ),ctrlOutVec( p )] = PixelStreamingDesignHDLDesign_fil( pixInVec( p )
%          end
```

Replace the commented loop with the code below. Call the step method of the fil object with
vectors containing the whole frame of data pixels and control signals. Pass each control signal to the
object separately, as a vector of logical values. Then, recombine the control signal vectors into a
vector of structures.

```
[pixOutVec,hStartOut,hEndOut,vStartOut,vEndOut,validOut] = ...
    fil(pixInVec,[ctrlInVec.hStart]',[ctrlInVec.hEnd]',[ctrlInVec.vStart]',[ctrlInVec.vEnd]',[ctr
ctrlOutVec = arrayfun(@(hStart,hEnd,vStart,vEnd,valid) ...
    struct('hStart',hStart,'hEnd',hEnd,'vStart',vStart,'vEnd',vEnd,'valid',valid),...
    hStartOut,hEndOut,vStartOut,vEndOut,validOut);
```

These code changes remove the pixel-by-pixel verification of the FIL results against the MATLAB results. Optionally, you can add a pixel loop to call the reference function, and a frame-by-frame comparison of the results. However, calling the original function for a reference slows down the simulation.

```
for p = 1:numPixPerFrm
    [ref_pixOutVec(p),ref_ctrlOutVec(p)] = PixelStreamingDesignHDLDesign(pixInVec(p),ctrlInVec(p
end
```

After the call to the `fil` object, compare the output vectors.

```
hdlverifier.assert(pixOutVec',ref_pixOutVec,'pixOut')
hdlverifier.assert([ctrlOutVec.hStart],[ref_ctrlOutVec.hStart],'hStart')
hdlverifier.assert([ctrlOutVec.hEnd],[ref_ctrlOutVec.hEnd],'hEnd')
hdlverifier.assert([ctrlOutVec.vStart],[ref_ctrlOutVec.vStart],'vStart')
hdlverifier.assert([ctrlOutVec.vEnd],[ref_ctrlOutVec.vEnd],'vEnc')
hdlverifier.assert([ctrlOutVec.valid],[ref_ctrlOutVec.valid],'valid')
```

This modified test bench sends an entire frame to the FPGA board in each packet, significantly improving the efficiency of the communication link.

See Also

Blocks
FIL Frame To Pixels | FIL Pixels To Frame | Image Filter

Objects
`visionhdl.ImageFilter`

More About

- "FPGA Verification" (HDL Verifier)

Prototype Vision Algorithms on Zynq-Based Hardware

You can use the Vision HDL Toolbox Support Package for Xilinx Zynq-Based Hardware to prototype your vision algorithms on Zynq-based hardware that is connected to real input and output video devices. Use the support package to:

- Capture input or output video from the board and import it into Simulink for algorithm development and verification.
- Generate and deploy vision IP cores to the FPGA on the board. (requires HDL Coder)
- Generate and deploy C code to the ARM processor on the board. You can route the video data from the FPGA into the ARM® processor to develop video processing algorithms targeted to the ARM processor. (requires Embedded Coder®)
- View the output of your algorithm on an HDMI device.

Video Capture

Using this support package, you can capture live video from your Zynq device and import it into Simulink. The video source can be an HDMI video input to the board, an on-chip test pattern generator included with the reference design, or the output of your custom algorithm on the board. You can select the color space and resolution of the input frames. The capture resolution must match that of your input camera.

Once you have video frames in Simulink, you can:

- Design frame-based video processing algorithms that operate on the live data input. Use blocks from the Computer Vision Toolbox™ libraries to quickly develop frame-based, floating-point algorithms.
- Use the Frame To Pixels block from Vision HDL Toolbox to convert the input to a pixel stream. Design and verify pixel-streaming algorithms using other blocks from the Vision HDL Toolbox libraries.

Reference Design

The Vision HDL Toolbox Support Package for Xilinx Zynq-Based Hardware provides a reference design for prototyping video algorithms on the Zynq boards.

When you generate an HDL IP core for your pixel-streaming design using HDL Workflow Advisor, the core is included in this reference design as the FPGA user logic section. Points **A** and **B** in the diagram show the options for capturing video into Simulink.

The FPGA user logic can also contain an optional interface to external frame buffer memory, which is not shown in the diagram.
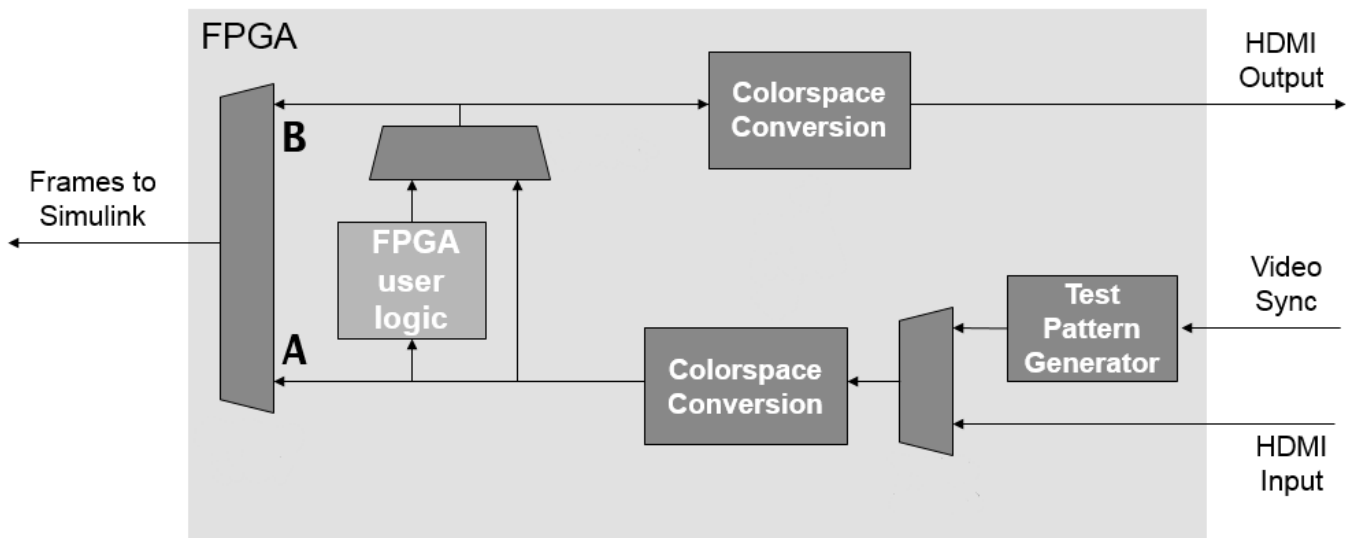
Note The reference design on the Zynq device requires the same video resolution and color format for the entire data path. The resolution you select must match that of your camera input. The design you target to the user logic section of the FPGA must not modify the frame size or color space of the video stream.

The reference design does not support multipixel streaming.

Deployment and Generated Models

By running all or part of your pixel-streaming design on the hardware, you speed up simulation of your video processing system and can verify its behavior on real hardware. To generate HDL code and deploy your design to the FPGA, you must have HDL Coder and the HDL Coder Support Package for Xilinx Zynq Platform, as well as Xilinx Vivado® and the Xilinx SDK.

After FPGA targeting, you can capture the live output frames from the FPGA user logic back to Simulink for further processing and analysis. You can also view the output on an HDMI output connected to your board. Using the generated hardware interface model, you can control the video capture options and read and write AXI-Lite ports on the FPGA user logic from Simulink during simulation.

The FPGA targeting step also generates a software interface model. This model supports software targeting to the Zynq hardware, including external mode, processor-in-the-loop, and full deployment. It provides data path control, and an interface to any AXI-Lite ports you defined on your FPGA targeted subsystem. From this model, you can generate ARM code that drives or responds to the AXI-Lite ports on the FPGA user logic. You can then deploy the code on the board to run along with the FPGA user logic. To deploy software to the ARM processor, you must have Embedded Coder and the Embedded Coder Support Package for Xilinx Zynq Platform.

See Also

More About

- "Vision HDL Toolbox Support Package for Xilinx Zynq-Based Hardware"

Block Reference Examples

Select Region of Interest

This example shows how to select a region of active frame from a video stream by using the ROI Selector block from the Vision HDL Toolbox™.

There are numerous applications where the input video is divided into several zones. In medical imaging, the boundaries of a tumor may be defined on an image or in a volume for the purpose of measuring its size. In geographical information systems (GIS), an ROI can be taken as a polygonal selection from a 2-D map.

Example Model
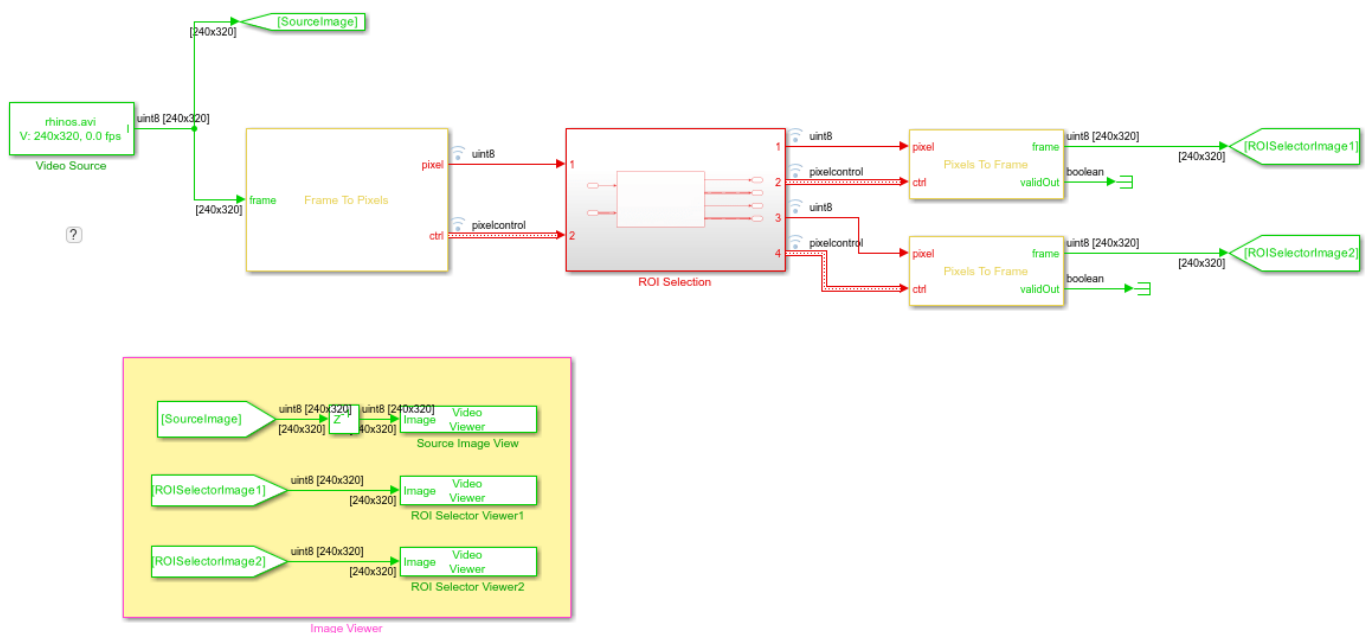
Copyright 2019 The MathWorks, Inc.

The example model includes a Video Source block that contains a 240p video sample. Each pixel is a scalar `uint8` value that represents intensity. The green and red lines represent full-frame processing and pixel-stream processing, respectively.

Serialize the Image

Use Frame To Pixels block to convert a full-frame image into pixel stream. To simulate the effect of horizontal and vertical blanking periods found in real life hardware video systems, the active image is augmented with non-image data. For more information on the streaming pixel protocol, see "Streaming Pixel Interface" on page 1-2. The Frame To Pixels block is configured as shown:
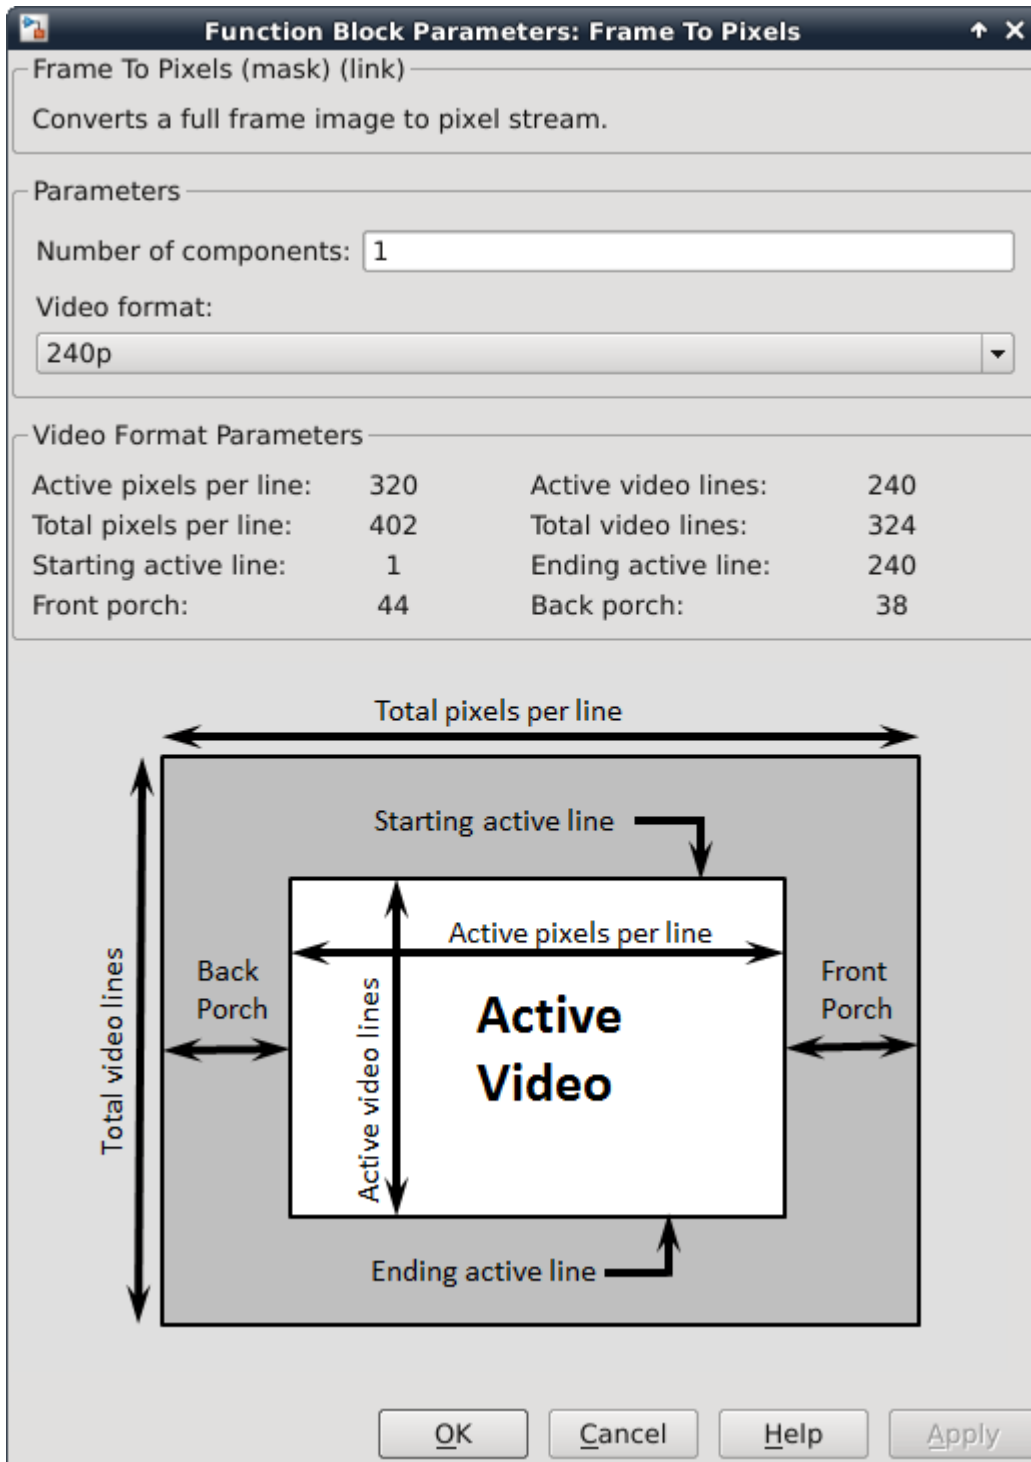
The **Number of components** parameter is set to 1 for grayscale image input, and the **Video format** parameter is 240p to match the video source.

In this example, the Active Video region corresponds to the 240x320 matrix of the source image. Six other parameters, namely, **Total pixels per line**, **Total video lines**, **Starting active line**, **Ending**

active line, **Front porch**, and **Back porch**, specify how many non-image data will be augmented on the four sides of the Active Video. For more information, see the Frame To Pixels block reference page.

Note that the sample time of the Video Source block is determined by the product of **Total pixels per line** and **Total video lines**.

Select Regions of Interest

The ROI Selection subsystem contains only an ROI Selector block.

Use the ROI Selector block to select regions of interest. You can use the **Regions** parameter to experiment with different region sizes and examine their effect on the output frames. In this model, the **Regions** parameter is set to `[100 100 50 50;220 170 100 70]` which represents two regions, each specified by `[hPos vPos hSize vSize]`. The first region is 50-by-50 pixels and located 100 pixels to the right and 100 pixels down from the top-left corner of the active frame. The second region is 100 pixels wide and 70 pixels tall, and is located in the bottom-right corner of the active frame.

The ROI Selector block accepts a pixel stream and a bus that contains five control signals from the Frame To Pixels block. It returns each region as a pixel stream that uses the same protocol, by manipulating the control signals. Each region is selected by setting the `valid` signal in the output `pixelcontrol` bus to `false` for any pixels not included in the requested region.
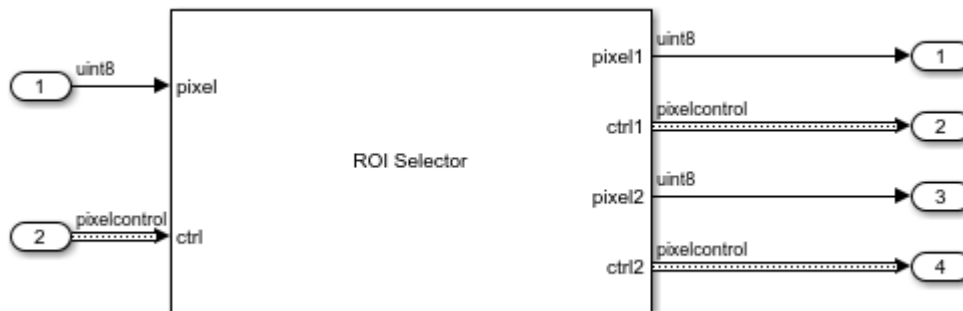
Display Regions of Interest

Use the Pixels To Frame block to convert the pixel stream back into a full frame. Since the output of the Pixels To Frame block is a 2-D matrix of a full image, there is no further need for the `pixelcontrol` bus.

The **Number of components** and **Video format** parameters of both Frame To Pixels and Pixels To Frame are set to `1` and `240p`, respectively, to match the format of the video source. The size of each active frame is smaller than 240p after the ROI selection. The Pixels to Frame block returns a 240-by-320 matrix with the active portion of the frame in the top-left corner.

Run the model to display the results. The model displays the output video streams by using three Video Viewer blocks.

- Source Image View -- The input video from the Video Source block
- ROI Selector Viewer1 -- The 50-by-50 pixel region
- ROI Selector Viewer2 -- The 100-by-70 pixel region

One frame of the source video and the two regions are shown from left to right.
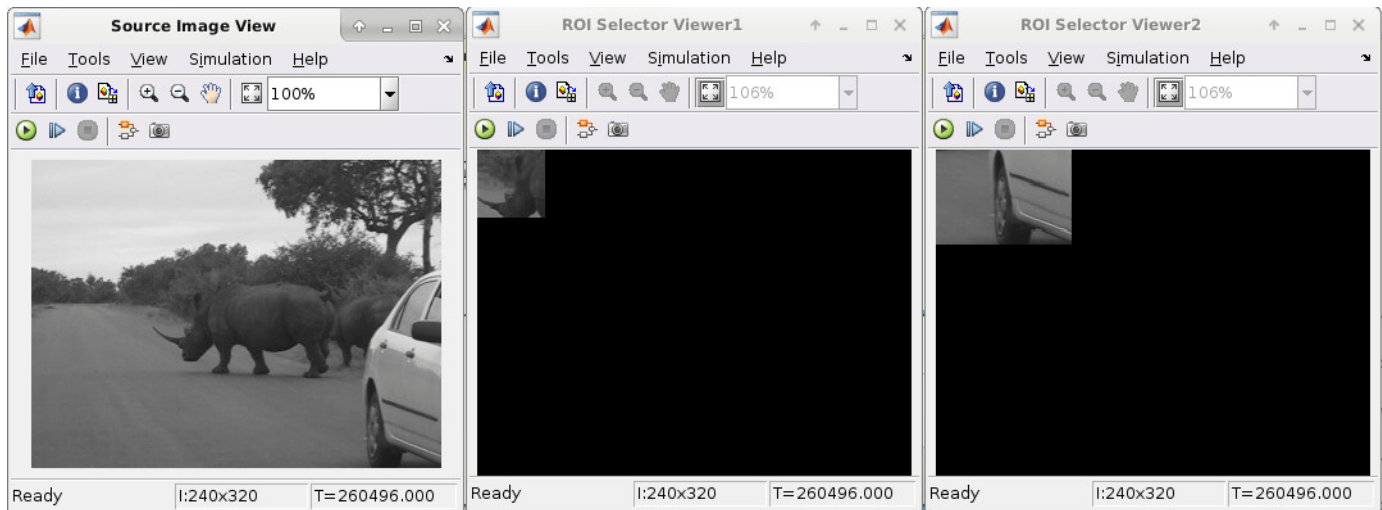
The Unit Delay block on the top level of the model is to time-align the matrices for a fair comparison.

Generate HDL Code

To check and generate the HDL code referenced in this example, you must have an HDL Coder™ license.

To generate the HDL code, use the following command:

```
makehdl('ROISelectionHDL/ROI Selection')
```

To generate a test bench, use the following command:

```
makehdltb('ROISelectionHDL/ROI Selection')
```

See Also

Blocks
Pixels To Frame | Frame To Pixels

Select Regions for Vertical Reuse

This example shows how to divide a frame into tiled regions of interest (ROIs) and use those regions to configure the ROI Selector block for vertical reuse.

Vertical reuse means dividing each frame into vertically-aligned regions where each column of regions shares a pixel stream. This arrangement enables parallel processing of each column, and the reuse of downstream processing logic for each region in the column.

Set up the size of the frame.

```
frmActiveLines = 240;
frmActivePixels = 320;
```

Divide the frame into equally-sized vertically-aligned regions, or tiles. The `visionhdlframetoregions` function returns an array of such regions, where each region is defined by four coordinates, and is of the form [*hPos vPos hSize vSize*]. These tile counts divide evenly into the frame dimensions, so no remainder pixels exist. The output regions cover the entire frame.

```
numHorTiles = 2;
numVerTiles = 2;
regions = visionhdlframetoregions(frmActivePixels,frmActiveLines,numHorTiles,numVerTiles)
```

```
regions =

     1     1   160   120
   161     1   160   120
     1   121   160   120
   161   121   160   120
```

The ROI Selector block in the Simulink model has the **Reuse output ports for vertically aligned regions** parameter selected, and uses the `regions` variable to define its output streams. The block has one output pixel stream per column of regions.

```
open_system('TiledROIHDL')
```
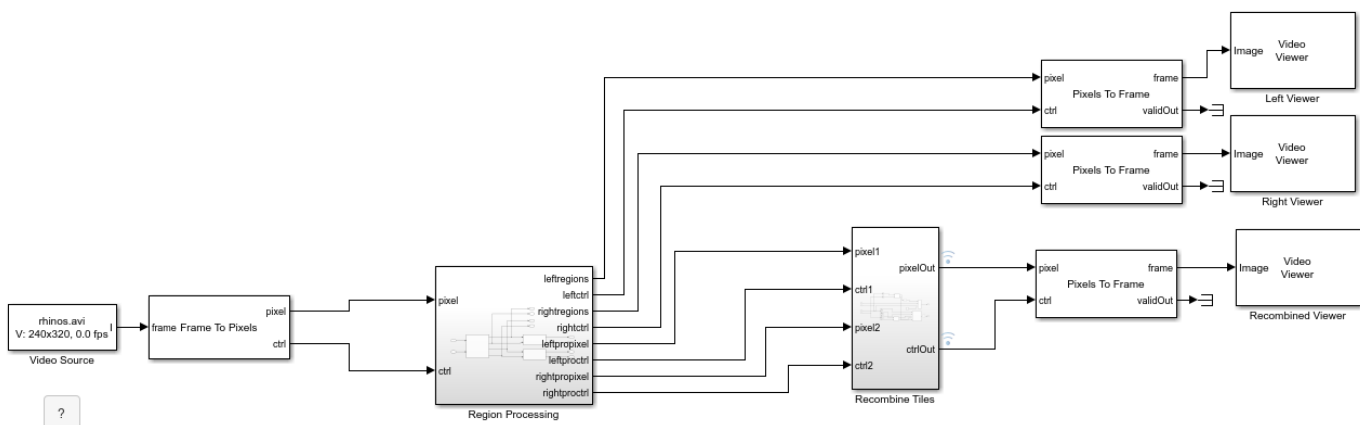

Copyright 2020 The MathWorks, Inc.

The start and end signals define each region in the pixel stream. When you run the model, you can see the output tiled regions changing in the Left Viewer and Right Viewer windows. The example performs opposite gamma correction operations on the left and right tiles, and then reassembles the four tiles into a complete frame by manipulating the `pixelcontrol` signals.

The blanking interval required by the downstream processing algorithm must be less than the interval between tiles. The blanking interval after each region is less than one line of pixels, so operations that require a vertical blanking interval, like those that use a line buffer, do not work. The gamma correction operation uses a lookup table that does not require a blanking interval.

```
sim('TiledROIHDL')
```
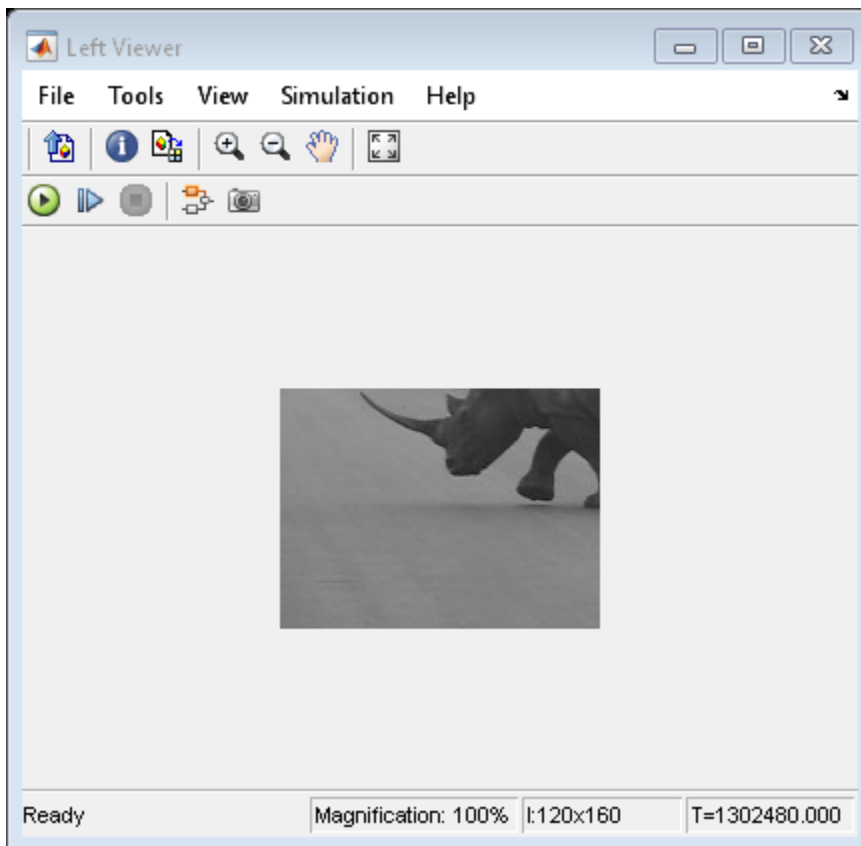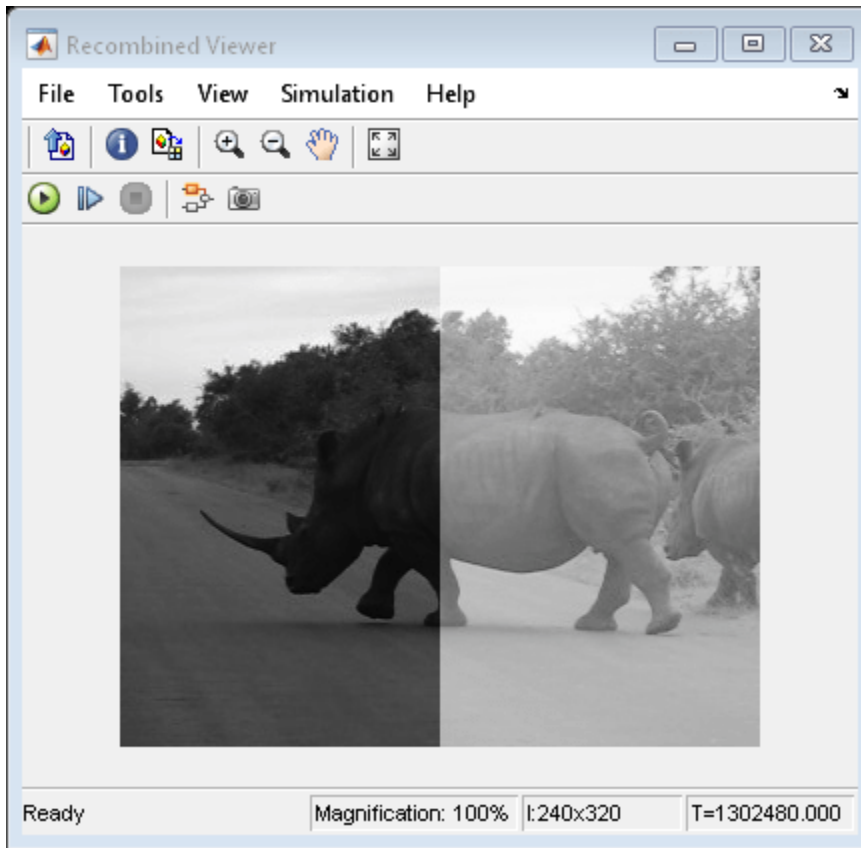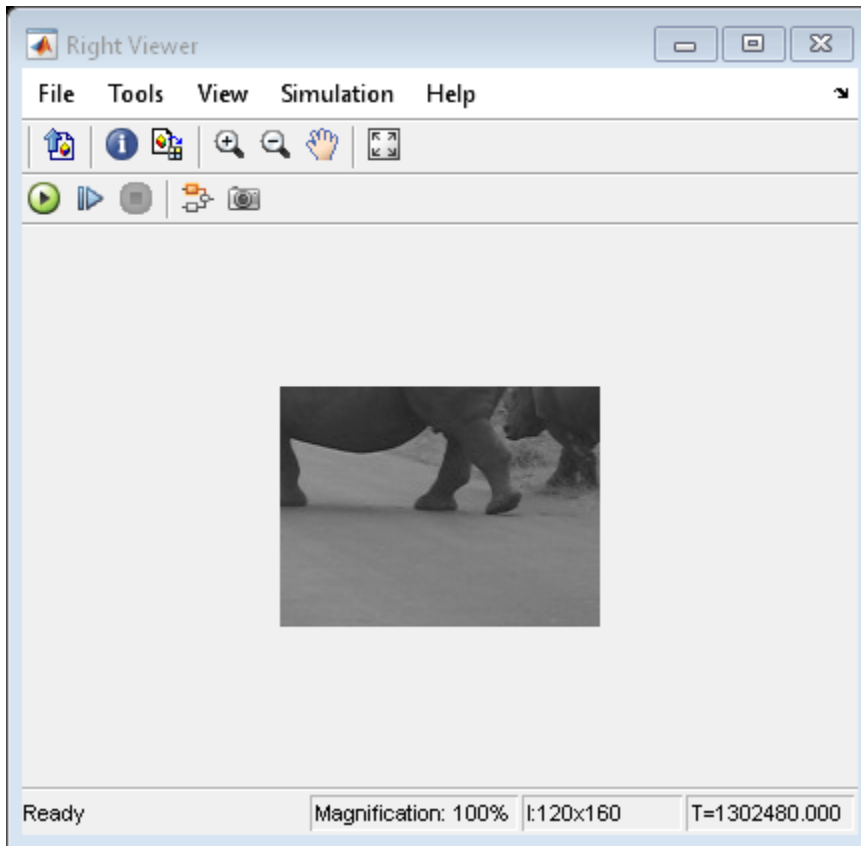
See Also

ROI Selector | `visionhdlframetoregions`

Construct a Filter Using Line Buffer

This example shows how to use the Line Buffer block to extract neighborhoods from an image for further processing. The model constructs a separable Gaussian filter.
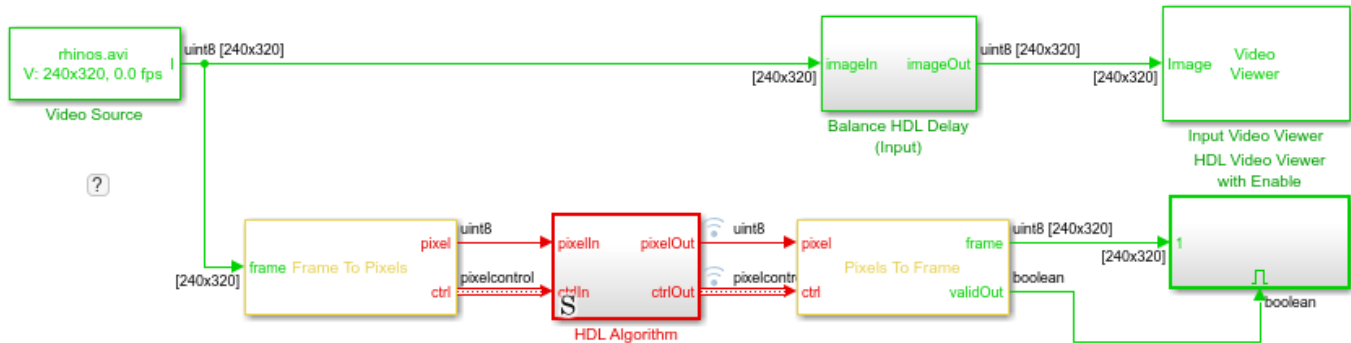
Copyright 2017 The MathWorks, Inc.

Inside the HDL Algorithm subsystem, the Line Buffer block is configured for a 5-by-5 neighborhood. The output is a 5-by-1 column vector. The Gain and Sum blocks implement separate horizontal and vertical components of a 5-by-5 Gaussian filter with a 0.75 standard deviation. After vertical filtering, the model stores the column sums in a shift register that creates a 1-by-5 row vector. The row values are filtered again to calculate the new central pixel value of each neighborhood.
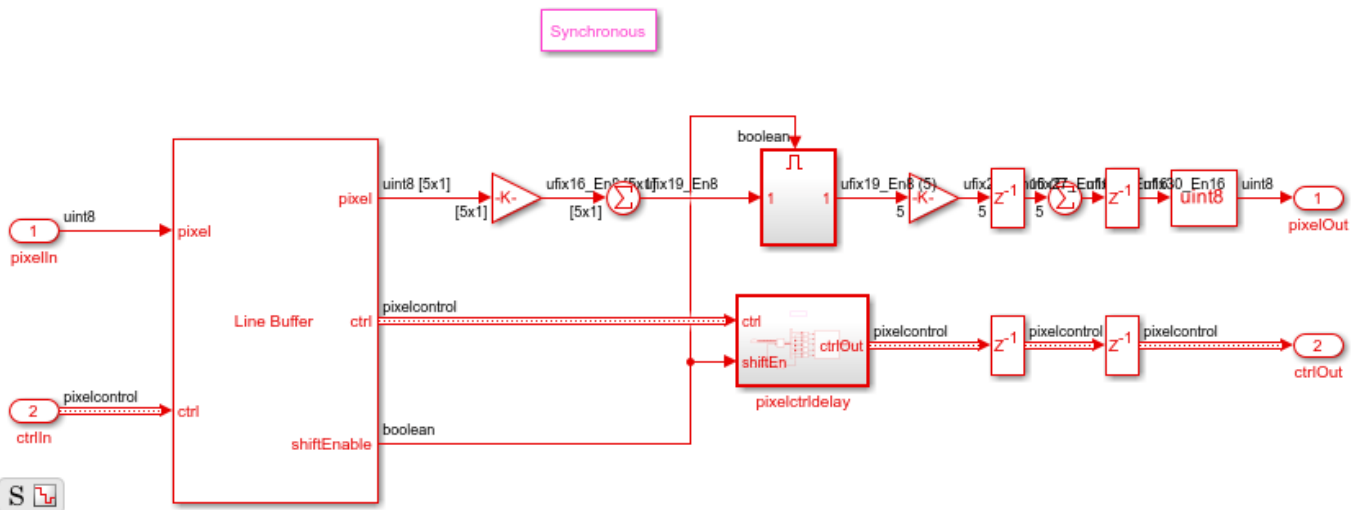
You can generate HDL code from the HDL Algorithm subsystem. You must have the HDL Coder™ software installed to run this command.

```
makehdl('SeparableFilterSimpleHDL/HDL Algorithm')
```

To generate an HDL test bench, use this command.

```
makehdltb('SeparableFilterSimpleHDL/HDL Algorithm')
```

See Also

Blocks
Frame To Pixels

Objects
visionhdl.LineBuffer

Convert RGB Image to YCbCr 4:2:2 Color Space

This example shows how to convert a pixel stream from R'G'B' color space to Y'CbCr 4:2:2 color space.
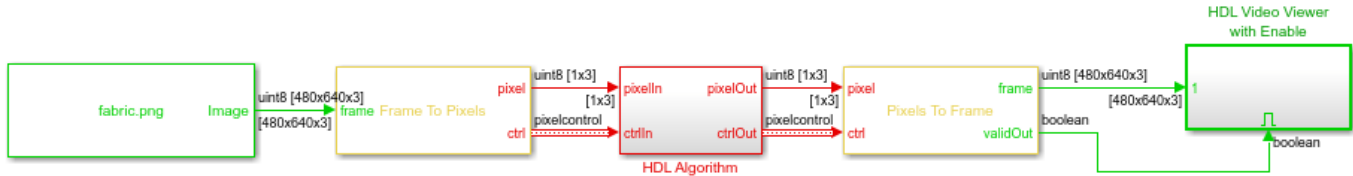
Copyright 2017-2022 The MathWorks, Inc.

The model imports a 480p RGB image, then the Frame to Pixels block converts it to a pixel stream. Inside the HDL Algorithm subsystem, the Color Space Converter and Chroma Resampler blocks convert the pixel stream to YCbCr 4:2:2 format.
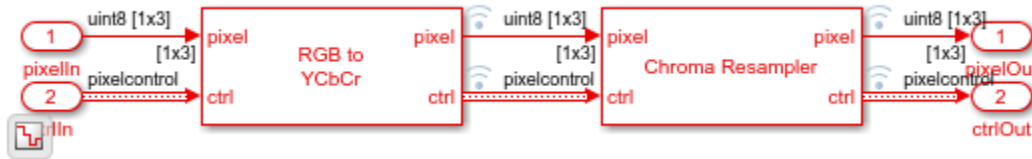
The waveform of the input and output pixel stream of the Chroma Resampler block shows the downsampling of the CbCr component values. The latency of the Chroma Resampler block depends on the size of the antialiasing filter. This example uses the default filter, which has 29 taps.
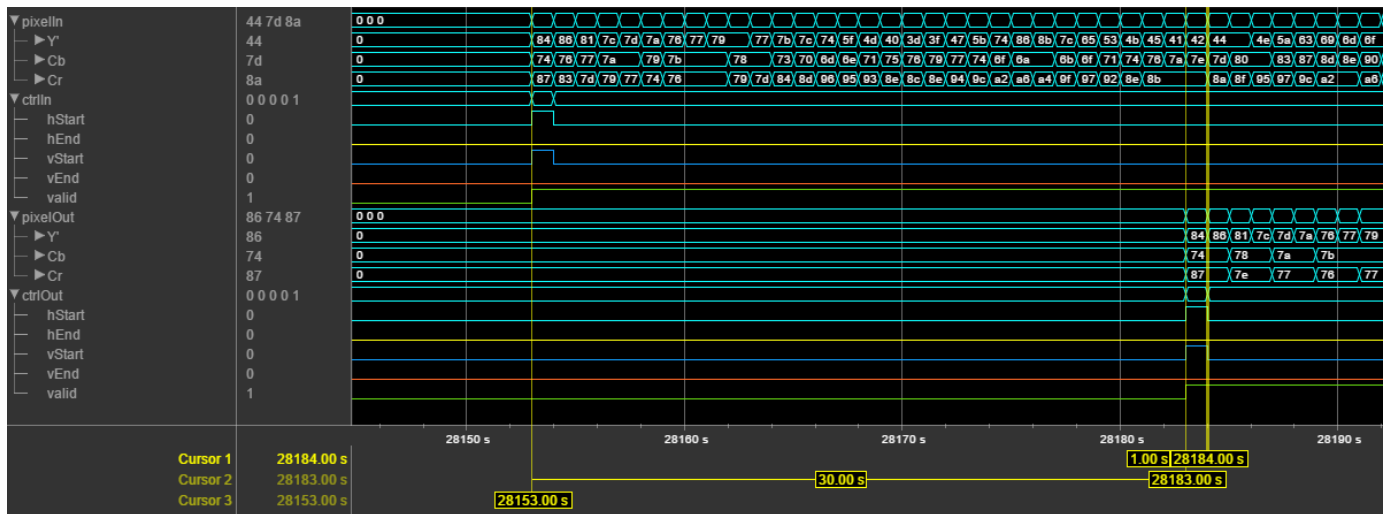
To check and generate the HDL code referenced in this example, you must have an HDL Coder™ license.

To generate the HDL code, use the following command.

```
makehdl('ChromaResampleExample/HDL Algorithm')
```

To generate the test bench, use the following command. Note that test bench generation takes a long time due to the large data size. Consider reducing the simulation time before generating the test bench.

```
makehdltb('ChromaResampleExample/HDL Algorithm')
```

The part of the model between the Frame to Pixels and Pixels to Frame blocks can be implemented on an FPGA.

See Also

Blocks
Frame To Pixels | Color Space Converter | Chroma Resampler

Compute Negative Image

Create the negative of an image by looking up the opposite pixel values in a table.
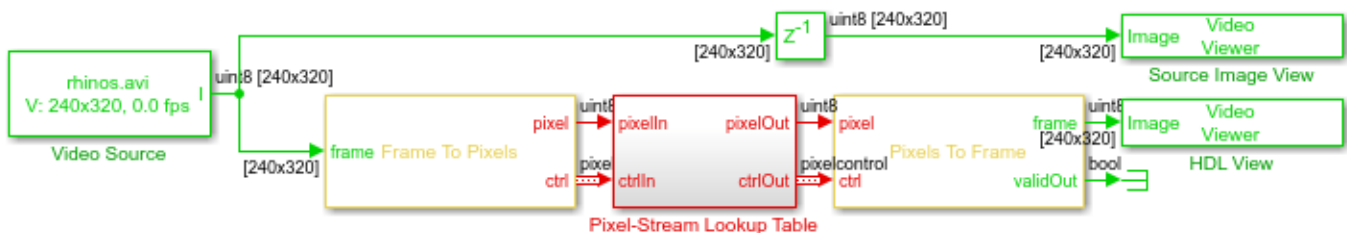
Copyright 2017 The MathWorks, Inc.

For a hardware-compatible design, the model converts the input video to a stream of pixel values. The Frame to Pixels and Pixels to Frame blocks are configured to match the format of the video source.

The Pixel-Stream Lookup Table subsystem contains a Lookup Table block, configured with inversion data. The input pixel data type is `uint8`, so the negative value is `255 - pixel`, or `linspace(255,0,256)`. The output pixel data type is the same as the data type of the table contents, in this case, `uint8`.

To generate and check the HDL code referenced in this example, you must have an HDL Coder™ license.

To generate the HDL code, use the following command:

```
makehdl('LookupTableHDL/Pixel-Stream Lookup Table')
```

To infer a RAM to implement the lookup table, the `LUTRegisterResetType` property is set to none. To access this property, right-click the **Lookup Table** block inside the subsystem, and navigate to **HDL Coder > HDL Block Properties**.

To generate a test bench for the generated HDL code, use the following command:

```
makehdltb('LookupTableHDL/Pixel-Stream Lookup Table')
```

See Also

Blocks
Frame To Pixels | Lookup Table

Adapt Image Filter Coefficients from Frame to Frame

This example shows how to use programmable coefficients to correct a time-varying impairment on the input video.

There are many different techniques for filtering image and video signals that require filter coefficients that vary from frame to frame. To dynamically change the coefficients of the Image Filter block, set the **Filter coefficients source** parameter to `Input port`. The Image Filter block samples the input coefficient port at the beginning of each frame.

The Example Model

The example model applies a brightness impairment to the input video, and the **HDL Filter** subsystem calculates filter coefficients for each frame and corrects the impairment. The model includes three video viewers: one for the original input video, another for the impaired video, and the third for the result of the filter that counteracts the impairment.

The model also includes Frame to Pixels and Pixels to Frame blocks to convert the matrix format video to streaming format suitable for HDL modeling.
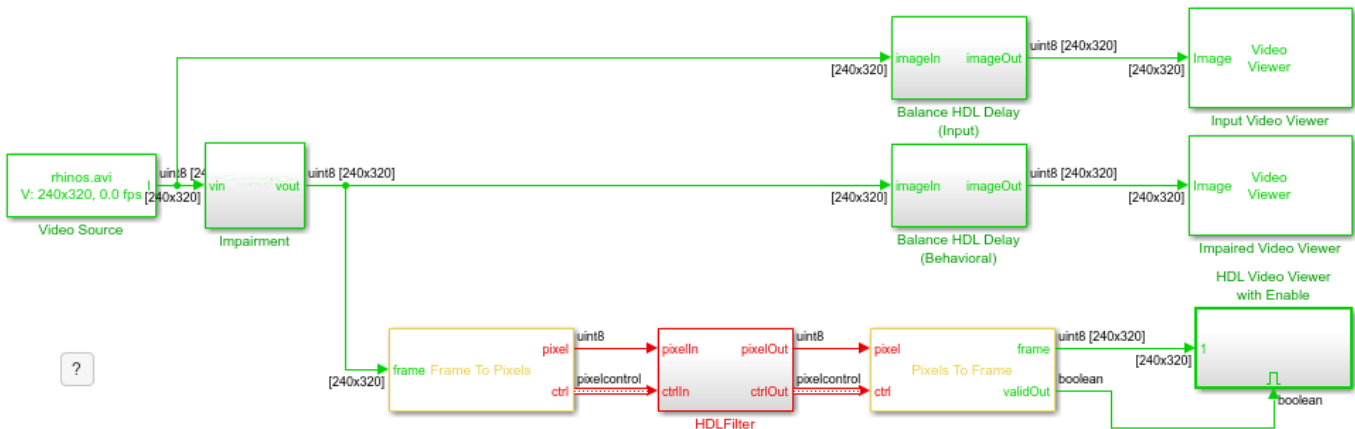
Copyright 2019 The MathWorks, Inc.

The Impairment

The impairment in this model is brightness modulation using a slow sine wave. Since the impairment is modeled purely behaviorally, the first step is to convert the image to double-precision values. The 16-bit counter counts up at the frame rate and the counter value is multiplied by 2*pi/40. The sine wave output is scaled down by 0.3 and a bias of 1.0 is then added. These calculations result in a +/-30% change in brightness over a period of 40 frames. After applying the impairment, the model converts back to `uint8` by using rounding with saturation.
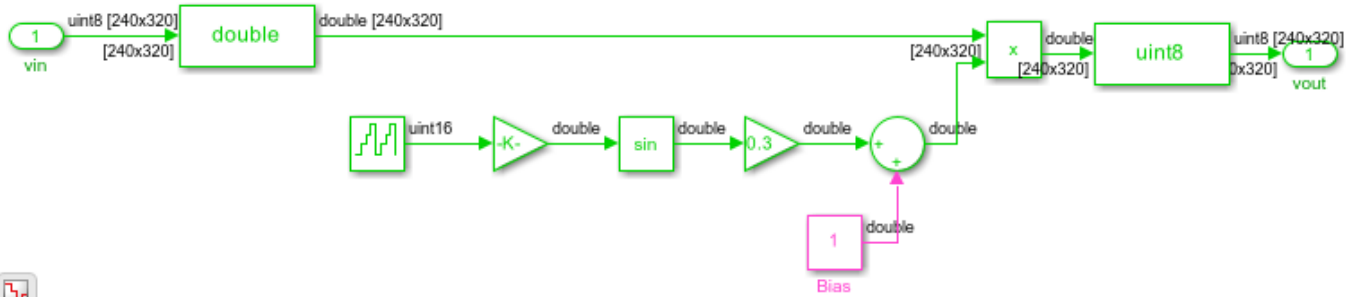
The Filter Algorithm

The HDL Algorithm subsystem starts by extracting a region of interest in the center of the image. Since this model is configured for a 320x240 video source, it uses a 100x100 region in the center of the video stream.

The Image Statistics block finds the mean of that central region. A new mean is computed for each 100x100 frame. The block sets the validOut port to `true` to indicate when the new mean is valid.
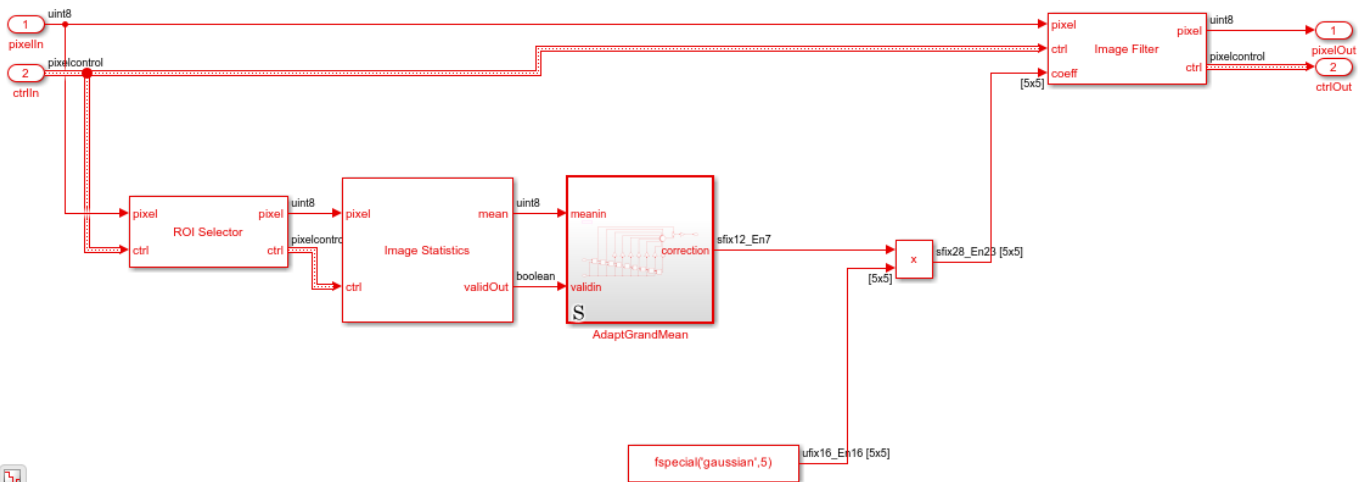
Compute the Scaled Grand Mean

The **Adapt Grand Mean** subsystem computes the correction factor required to counteract the impairment.

You could use central-region-mean brightness directly, with a "gray-world" assumption that the average brightness is mid-scale (128 in this case). But, a more accurate approach is to use the previous brightness means, with the assumption that the average brightness does not change quickly frame to frame.

Forming a mean of means is known as a grand mean, but that calculation would give equal weight to the past frames. Instead, the subsystem weights the past frames with an exponential fractional decay with the coefficients [1 1/2 1/4 1/8 1/16 1/32 1/64 1/64]. The last coefficient would normally be 1/128 but by adjusting that value, the sum of the weights becomes exactly 2, making the normalizing factor a simple shift operation. Note that the initial value of all the delay line registers is mid-scale (128) to avoid large start-up transients in the correction.

The subsystem finds the correction factor using the current mean and the weighted grand mean. Since the grand mean scaled up by 2, if you subtract the current mean from it, the resulting value is the weighted grand mean plus or minus the error term in the direction of correcting the error.

The correction is then scaled by 2^-7 and sent to the output port. A normalization could be applied here by dividing by the grand mean, but in practice, simple scaling works well enough.
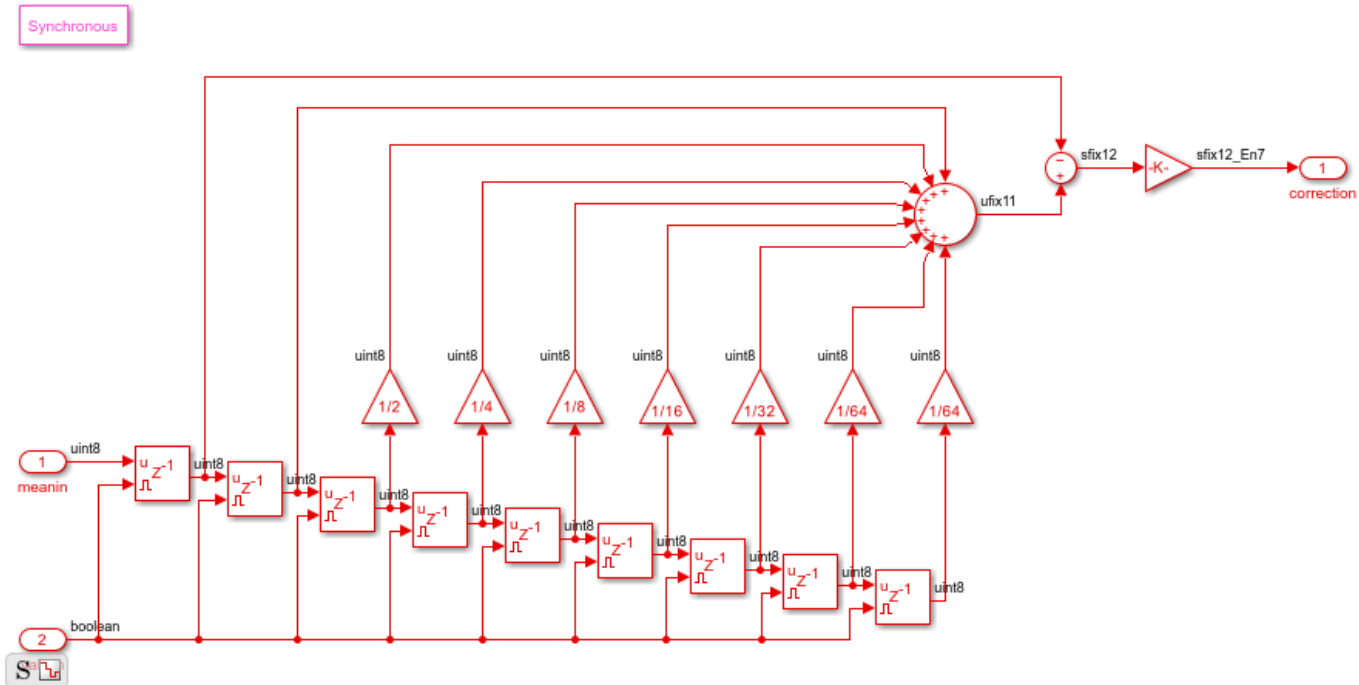
Apply the Correction

The correction output from the **Adapt Grand Mean** subsystem is then used to scale the filter coefficients, in this case a Gaussian filter of size 5x5 with a standard deviation of the default 0.5. In the actual FPGA this filter uses 25 multipliers. Pipelining is of no concern here since these values are computed well before they are needed. The block samples the coefficient port when the vStart signal in the input ctrl bus is true.

Going Further

In this simple example, you could alternatively apply the correction factor to the scalar pixel stream and then filter. The architecture shown can expand for more complex adaptive changes in the filter coefficients.

The 5x5 multiply of the correction factor with the gaussian coefficients could be implemented as a single serial multiplier rather than 25 parallel multipliers. To achieve this HDL implementation, include the Product block in a Subsystem, and right-click the Subsystem to open the HDL Block Properties. Set the **SharingFactor** property to 25 to implement a single time-multiplexed multiplier.

With this setting, the multiply operation uses a 25-times faster clock than the rest of the design. Consider your required pixel clock speed and whether your device can accommodate the faster rate.

See Also

Blocks
Image Filter | ROI Selector | Image Statistics

Video Stabilization

This example shows how to implement a feature-based video stabilization algorithm for FPGAs.

This algorithm helps reduce shaking between frames. Digital video stabilization techniques provide a more feasible and economical solution than a physical stabilizer.

This example implements the same algorithm as the "Video Stabilization Using Point Feature Matching" (Computer Vision Toolbox) example. This video stabilization algorithm comprises these steps.

1   Features from accelerated segment test (FAST) feature detector — Detect corners as features to match between the two frames.
2   Binary robust independent elementary features (BRIEF) descriptor — Calculate a unique description for each feature.
3   Feature matching — Match each feature description with its corresponding feature in the other frame .
4   Random sample consensus (RANSAC) affine transformation estimation — Calculate a transform that describes the movement of the second frame relative to the first one.

The `videoStabilization` subsystem has four output ports.

- **ransac_done** indicates the completion of RANSAC.
- **matrix_H** is the transform matrix in the form [[a, d, 0], [b, e, 0], [c, f, 1]].
- **orb_done** indicates the ORB (oriented FAST rotated BRIEF) part is done.
- **matched_points** is the matched points from ORB matching.

The last two output signals are for additional information or debugging.

```
modelname = 'VideoStabilizationHDL';
open_system(modelname);
set_param(modelname,'SampleTimeColors','on');
set_param(modelname,'Open','on');
set(allchild(0),'Visible','off');
```
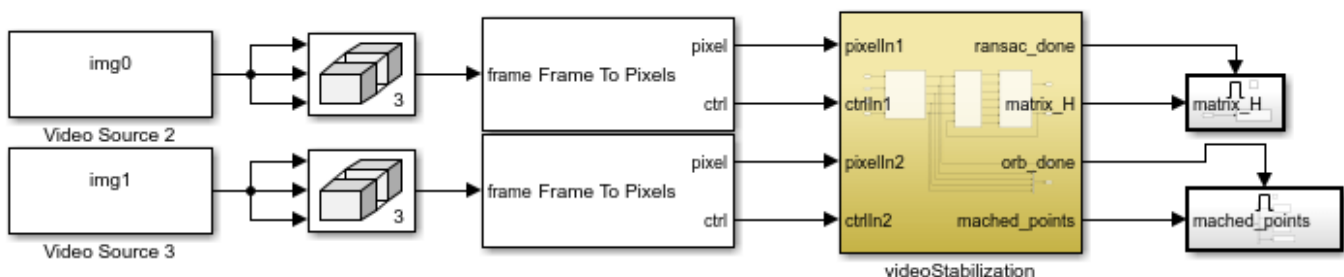

Copyright 2021 The MathWorks, Inc.

These images show consecutive frames from a camera.

4-19

```
imgA = imread('VideoStabilizationHDLExample_img_0.png');
imgB = imread('VideoStabilizationHDLExample_img_1.png');
figure; imshowpair(imgA, imgB, 'montage');
title(['Frame A', repmat(' ',[1 70]), 'Frame B']);
```


To highlight the difference, the figure shows a red-cyan color composite to show the pixel-wise difference between them. Frame A is shown in red, and frame B is shown in cyan.

```
figure; imshowpair(imgA,imgB,'ColorChannels','red-cyan');
```

FAST Feature Detection

This part of the algorithm follows the design concept outlined in the "FAST Corner Detection" on page 2-53 example to design a 9-16 FAST detector. The 9-16 FAST feature point detector finds the pixel value difference between the central pixel and its surrounding pixels, shown in this figure.

1	8	15	22	29	36	43
2	9	16	23	30	37	44
3	10	17	24	31	38	45
4	11	18	25	32	39	46
5	12	19	26	33	40	47
6	13	20	27	34	41	48
7	14	21	28	35	42	49

If seven of 12 pixels have a larger pixel value difference than a specified threshold, the detector considers the central point as a feature point. Then, the `NonMaxSuppress` subsystem filters out features that have less than the maximum metric within a 5-by-5 patch window. The figure shows the FAST subsystem.

```
open_system([modelname '/videoStabilization/oFAST_rBRIEF/FAST_1'],'force');
```
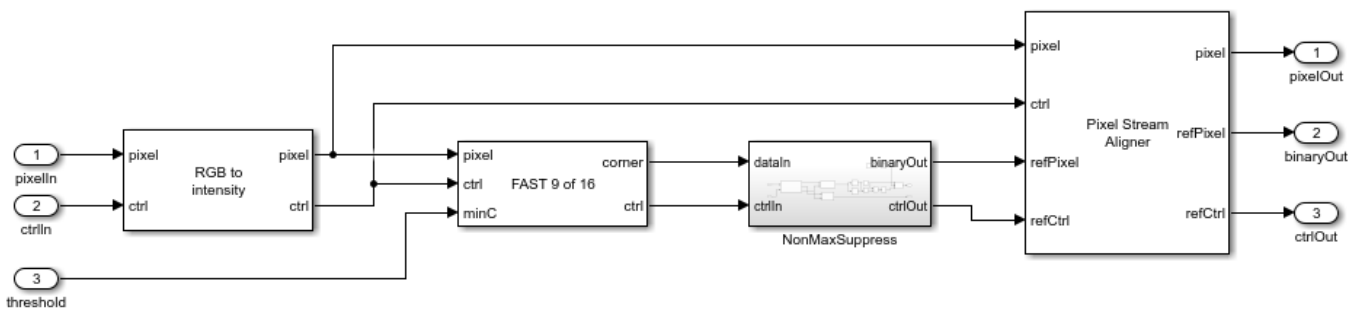

This figure shows the detected corners.

```
imgA_fast = imread('VideoStabilizationHDLExample_a_fast.png');
imgB_fast = imread('VideoStabilizationHDLExample_b_fast.png');
figure; imshowpair(imgA_fast, imgB_fast, 'montage');
title([ 'Corners in Frame A', repmat(' ',[1 40]),'Corners in Frame B']);
```

Corners in Frame A   Corners in Frame B
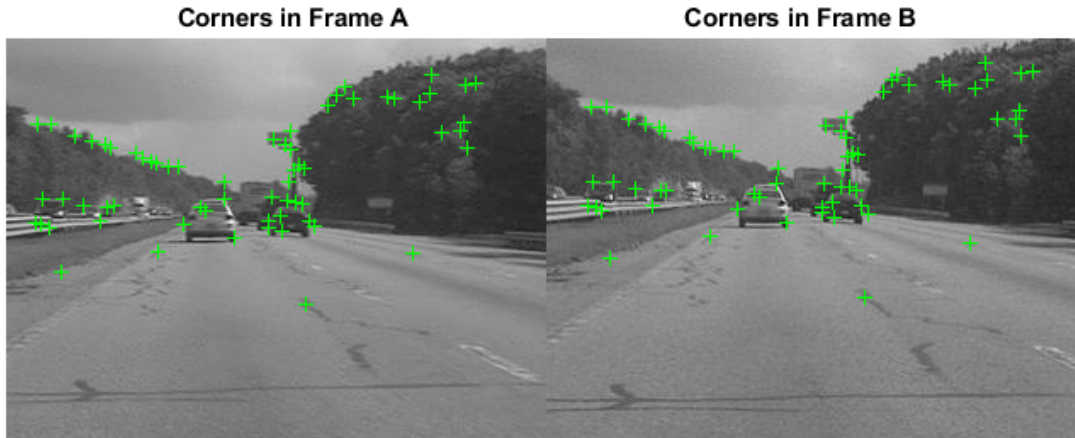
BRIEF Descriptor Generation

The purpose of a feature descriptor is to generate a unique string for each feature point. The algorithm uses these strings to match points between frames. The BRIEF descriptor generates a binary vector for each feature point returned from corner detection by comparing the pixel values in a fixed pattern [3]. This example uses oriented FAST features (oFAST) to increase robustness to image rotation. The oFAST algorithm rotates each image patch to align it to a common direction before generating its binary vector. The angle of rotation is the angle between the intensity centroids of both *x*- and *y*-axis, as defined by this equation.

$$m_{pq} = \sum_{x,y} x^p y^q I(x, y)$$

The algorithm calculates the intensity centroid by using a circular patch. In this case, the radius of the circle is 11. The algorithm calculates the weighted sum of a square patch and then subtracts the values at the corners. This figure shows the creation of a circular patch from a square patch.
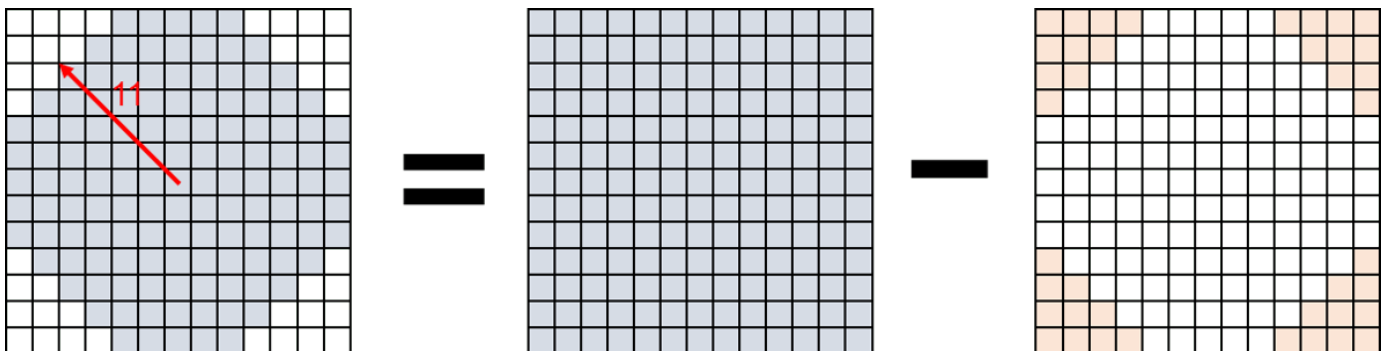
After intensity centroid calculation, the Complex to Magnitude-Angle block computes the rotation angle. The figure shows the BRIEF descriptor generator.

```
open_system([modelname '/videoStabilization/oFAST_rBRIEF/BRIEF_DESC_GEN_1/IntensityCentroid'],'f
```
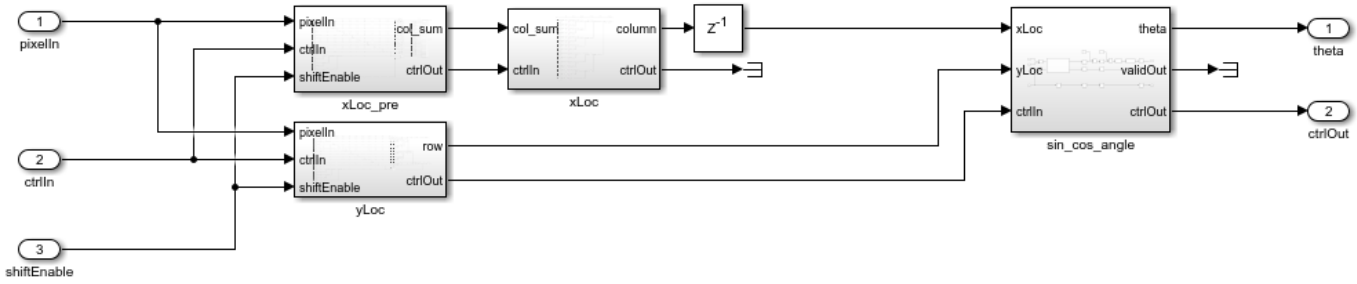
Before applying the fixed pattern, the BRIEF block rotates this circular patch by the computed angle. The fixed pattern used in this model comes from OpenCV software [3].

```
open_system([modelname '/videoStabilization/oFAST_rBRIEF/BRIEF_DESC_GEN_1/descriptor_gen'],'force
```
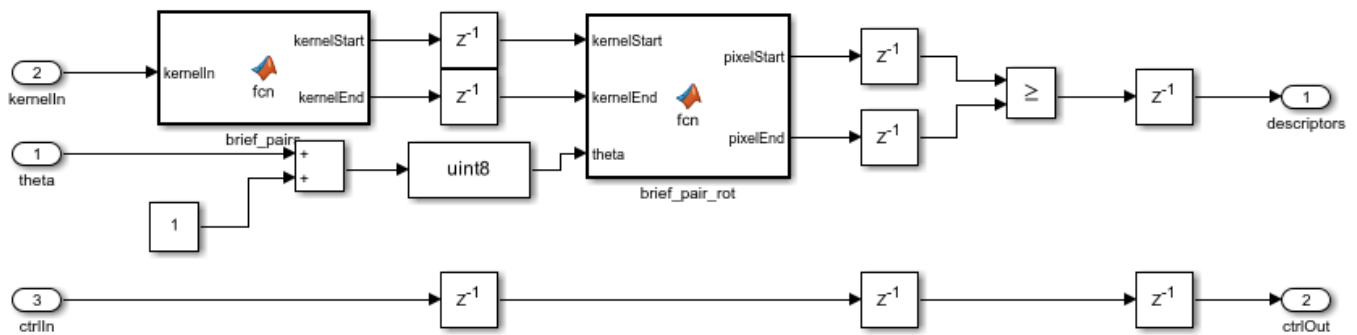

Because the BRIEF algorithm compares the pixel value of each fixed pair, the description generator is an array of comparators. The model stores the output binary vectors in the on-chip memory. The example supports 1024 features by default. You can change the number of features by setting the **Estimated number of feature points** parameter of the oFAST_rBRIEF subsystem. This parameter configures the size of the Simple Dual Port RAM blocks in the BRIEF_DESC_GEN_1 and descriptor_memory subsystems. This figure shows the BRIEF descriptor generation subsystem.

```
open_system([modelname '/videoStabilization/oFAST_rBRIEF/BRIEF_DESC_GEN_1'],'force');
```
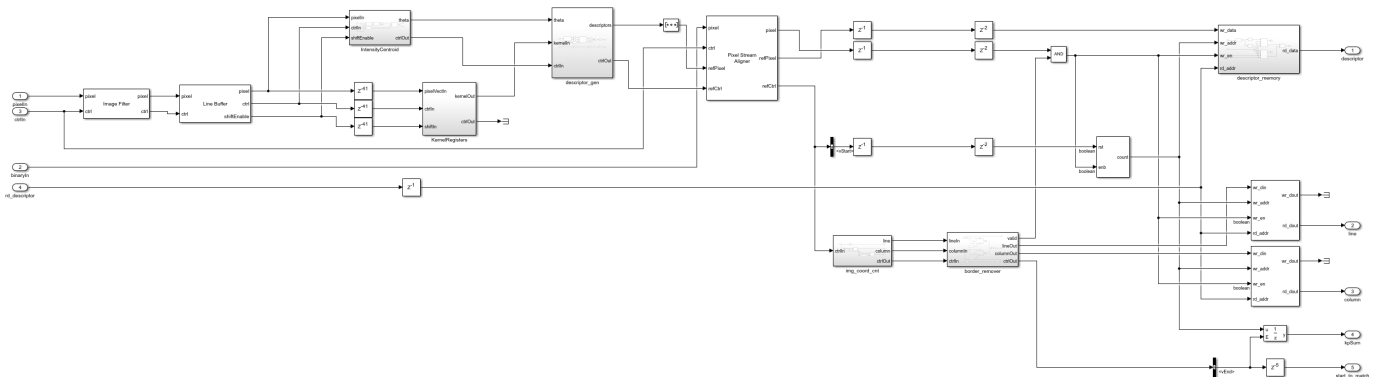

Feature Matching

This figure shows the matching process that uses the binary BRIEF feature descriptors.
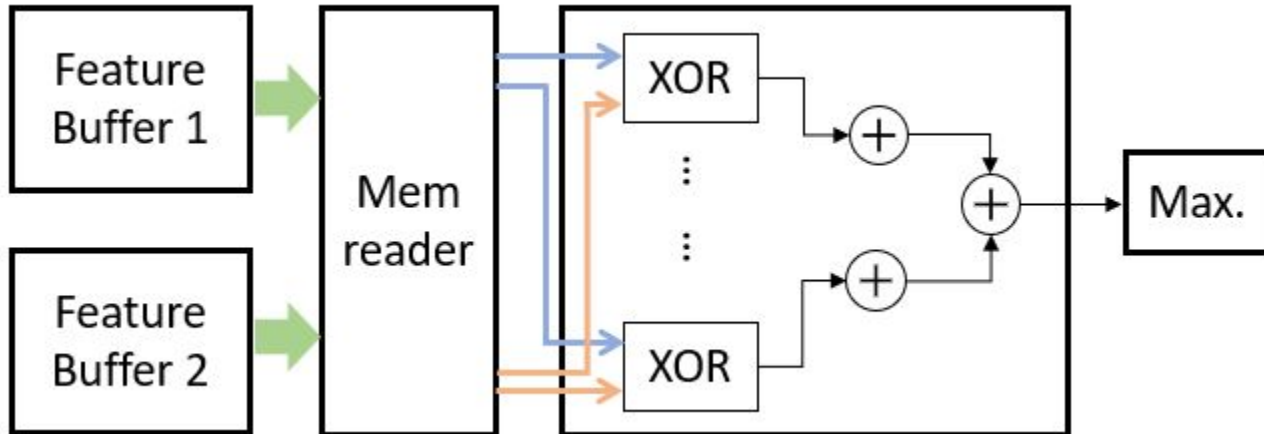
Because of the complexity of implementing sorting algorithms on hardware, this example implements a simple bubble-sorting algorithm to find the 64 best matched point pairs from two images. The bubble-sorting algorithm consumes fewer resources than a parallel sort. Because each image typically has hundreds of feature points, finding each pair of points would cost hundreds of clock cycles if you performed exhaustive matching. The time used by bubble sorting is much shorter than an exhaustive match.

```
open_system([modelname '/videoStabilization/oFAST_rBRIEF/BRIEF_MATCH'],'force');
```
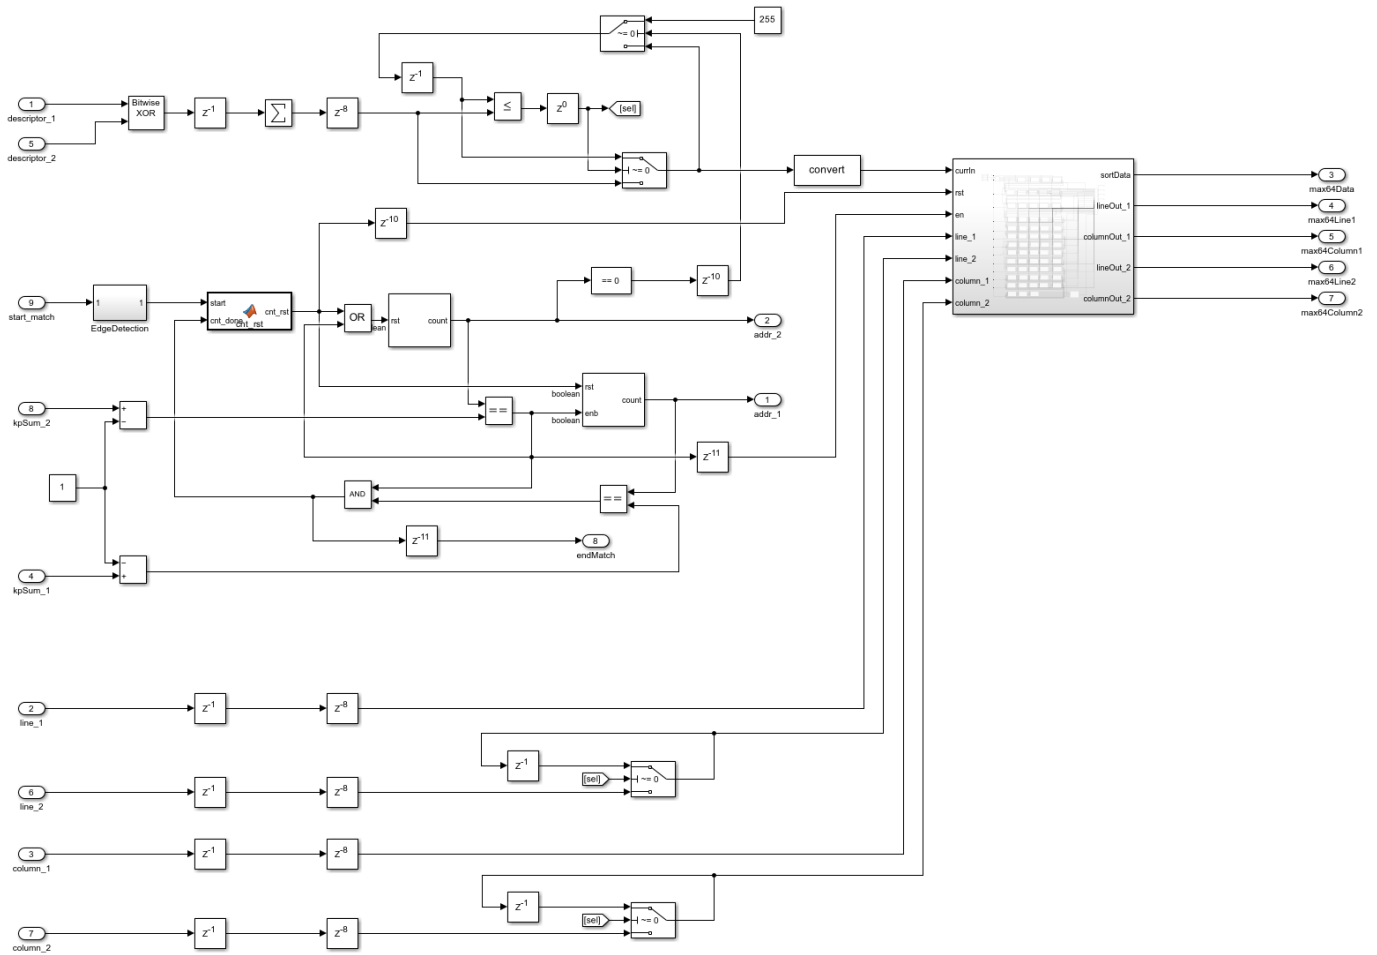
This image shows the results of matching feature points.
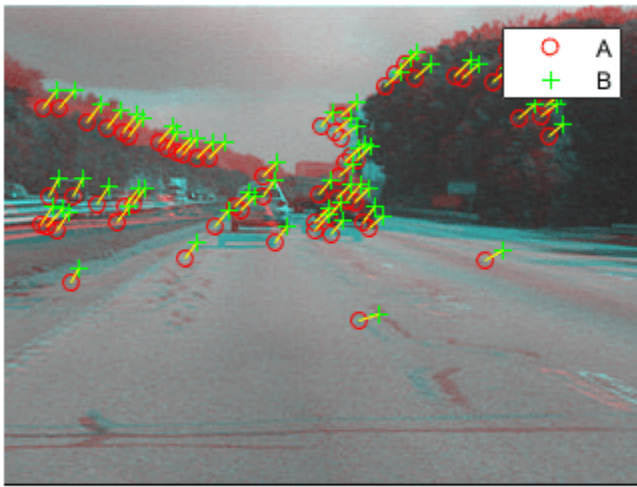
RANSAC Affine Transformation Estimation

RANSAC is a common algorithm for finding an optimal affine transform. The key idea of RANSAC is to repeatedly pick three random point pairs, and then calculate the affine transform matrix. The algorithm stops repeating when either 95% of the point pairs match (error is less than 2 pixels) or the algorithm completes the maximum number of iterations. You can change these parameters in the RANSAC subsystem mask.

```
open_system([modelname '/videoStabilization/RANSAC'],'force');
```
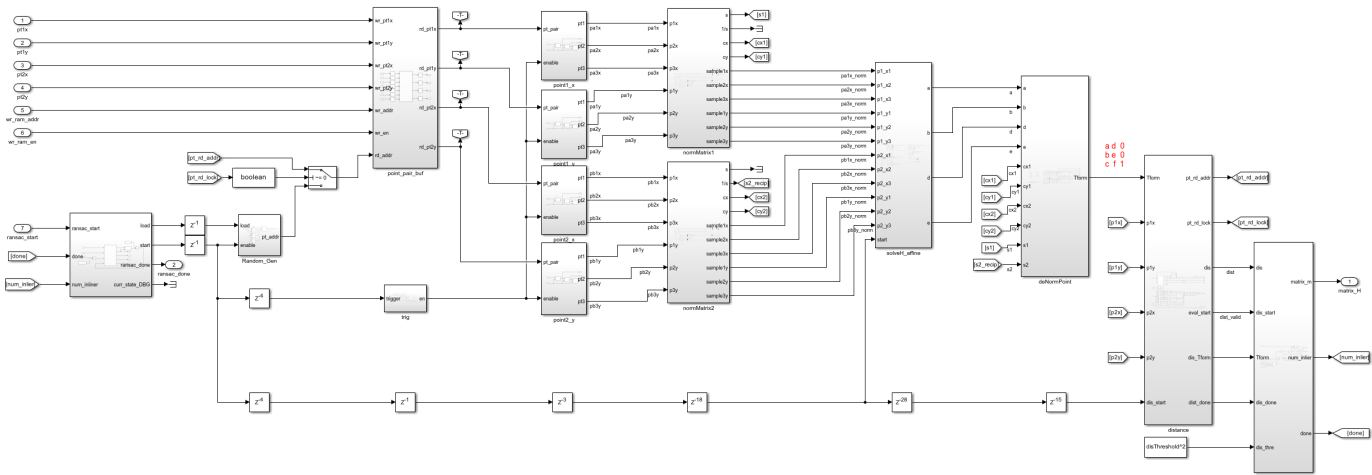

The standard procedure for RANSAC calculation comprises these steps.

1   Random selection — To reduce the computation burden of RANSAC, give the algorithm only the best matched 64 point pairs. The linear-feedback shift register (LFSR) randomly selects 3 of these 64 point pairs.

2   Normalization — Normalize point location *(x,y)* into 1.x format. This normalization is critical to guarantee a stable result from RANSAC.

3   Calculate affine matrix — Solve linear equations with six variables. This block implements the numerical solution of these equations. Because additional latency at this step does not affect the throughput of the design, this subsystem implements a fully pipelined simple divider. The simple divider uses fewer resources than other divider implementations for hardware. #Denormalization — Convert the point location back to *(x,y)* format. This operation is the inverse of normalization.

Model Configuration

You can configure these parameters on the subsystems.

Nonmaximal suppression — To reduce processing of extra features, enable nonmaximal suppression. You can also select the patch size. Typical patch sizes are five or seven pixels. Larger patch sizes result in fewer features with low FAST scores, but longer latency.

Minimum contrast for FAST — The typical range for this threshold is 15 to 20, but the value varies for different applications.

Estimated number of feature points — This parameter determines the depth of on-chip memory. This parameter is defined by this equation.

$2^{ceil(\log_2 N)}$, where N is the estimated number of features.

For example, the default value is `1000`, which results in an on-chip memory depth of 1024 to store feature descriptors.
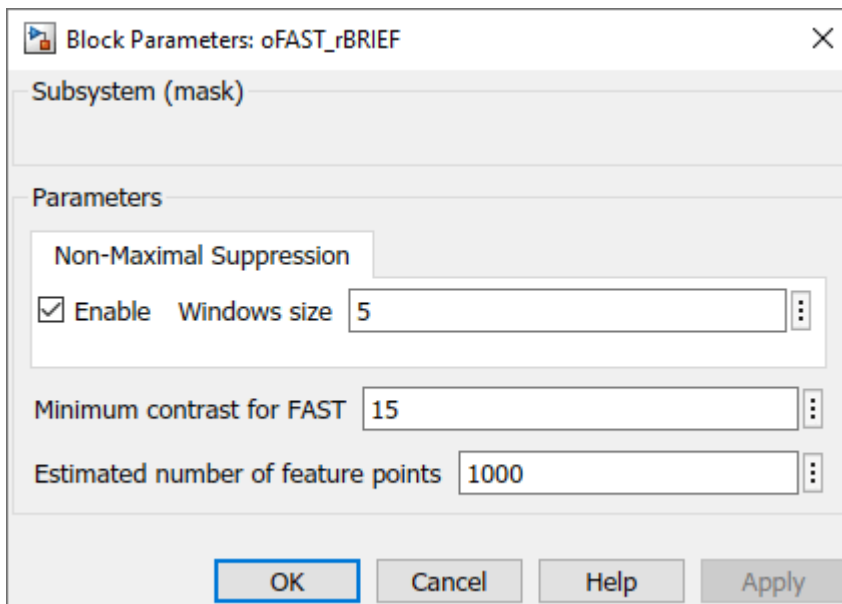
Maximum random trials — The default number of iterations is `1000`. This value is usually less than 2000.

Confidence of finding maximum number of inliers — Percentage of point pairs matched. Because this algorithm uses only the 64 best-matched pairs, this example uses a confidence value of 95%. If your design uses more point pairs, you can reduce this parameter.

Maximum distance from point to projection — The recommended range of this parameter is from 1.5 to 2.
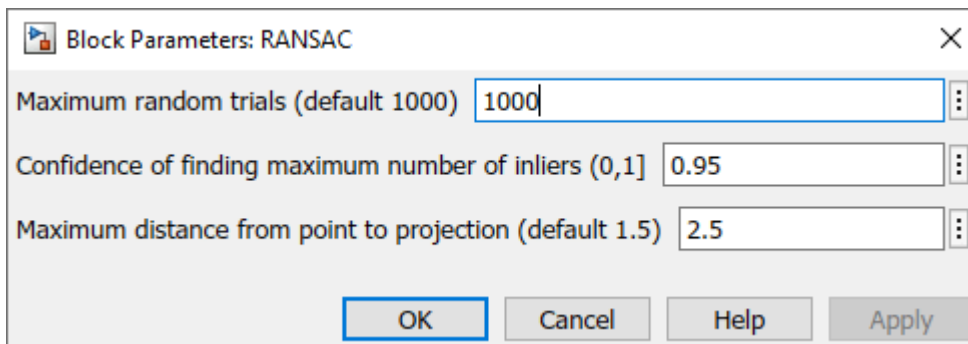
Performance

The figure shows the stabilization result obtained by applying the generated affine transform to the second input frame.

Because the camera is mounted on a moving vehicle, it has a different relative speed to the ground and to other vehicles on the road. Therefore, the matched point varies according to its position on the image. One way to extend this algorithm for applications with very large object size changes is to use an image pyramid to find features at different scales.

This design takes about 5 minutes to run an HDL simulation for 240p image/video with no more than 1024 feature points. This table shows the resource consumption on the Xilinx® Zynq®-7000 SoC ZC706 development kit.

Name	Slice LUTs (218600)	Slice Registers (437200)	F7 Muxes (109300)	F8 Muxes (54650)	Slice (54650)	LUT as Logic (218600)	LUT as Memory (70400)	Block RAM Tile (545)	DSPs (900)
∨ N videoStabilization	87988	84612	14288	150	30829	83645	4343	71	118
∨ □ u_oFAST_rBRIEF (oFAST_rBRIEF)	82846	77485	13908	0	29268	78711	4135	69	60
> □ u_BRIEF_DESC_GEN_1 (BRIEF_D	36381	30753	7002	0	13029	34429	1952	28.5	27
> □ u_BRIEF_DESC_GEN_2 (BRIEF_D	36301	30754	6906	0	12354	34344	1957	28.5	27
> □ u_BRIEF_MATCH (BRIEF_MATCH)	6227	8726	0	0	2400	6200	27	0	0
> □ u_FAST_1 (FAST_1)	1968	3626	0	0	912	1868	100	6	3
> □ u_FAST_2 (FAST_2)	1969	3626	0	0	857	1870	99	6	3
> □ u_point_pair_load (point_pair_load)	31	2170	0	0	379	13	18	0	0
> □ u_RANSAC (RANSAC)	4849	4957	380	150	1643	4659	190	2	58

References

[1] Rosten, E., and T. Drummond. "Fusing Points and Lines for High Performance Tracking." *Proceedings of the IEEE International Conference on Computer Vision* 2 (October 2005): 1508-1511.

[2] Rosten, E., and T. Drummond. "Machine Learning for High-Speed Corner Detection." *Computer Vision - ECCV 2006 Lecture Notes in Computer Science* (2006): 430-43. doi:10.1007/11744023_34.

[3] "OpenCV: Open Source Computer Vision Library" https://github.com/opencv/opencv, 2017

[4] Rublee, E., Rabaud, V., Konolige, K., and Bradski, G.. "ORB: An efficient alternative to SIFT or SURF." *Proceedings of the IEEE International Conference on Computer Vision* (2011): 2564-2571.

[5] Fischler, M.A., and Bolles, R.C.. "Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography." *Communications of the ACM* 24(6): 381-395.

5

Simulation Data Inspector

View Data in the Simulation Data Inspector

You can use the Simulation Data Inspector to visualize the data you generate throughout the design process. Simulation data that you log in a Simulink model logs to the Simulation Data inspector. You can also import test data and other recorded data into the Simulation Data Inspector to inspect and analyze it alongside the logged simulation data. The Simulation Data Inspector offers several types of plots, which allow you to easily create complex visualizations of your data.

View Logged Data

Logged signals as well as outputs and states logged using the `Dataset` format automatically log to the Simulation Data Inspector when you simulate a model. You can also record other kinds of simulation data so the data appears in the Simulation Data Inspector at the end of the simulation. To see states and output data logged using a format other than `Dataset` in the Simulation Data Inspector, open the Configuration Parameters dialog box and, in the **Data Import/Export** pane, select the **Record logged workspace data in Simulation Data Inspector** parameter.

Note When you log states and outputs using the `Structure` or `Array` format, you must also log time for the data to record to the Simulation Data Inspector.

The Simulation Data Inspector displays available data in the table in the **Inspect** pane. To plot a signal, select the check box next to the signal. You can modify the layout and add different visualizations to analyze the simulation data. For more information, see "Create Plots Using the Simulation Data Inspector" (Simulink).
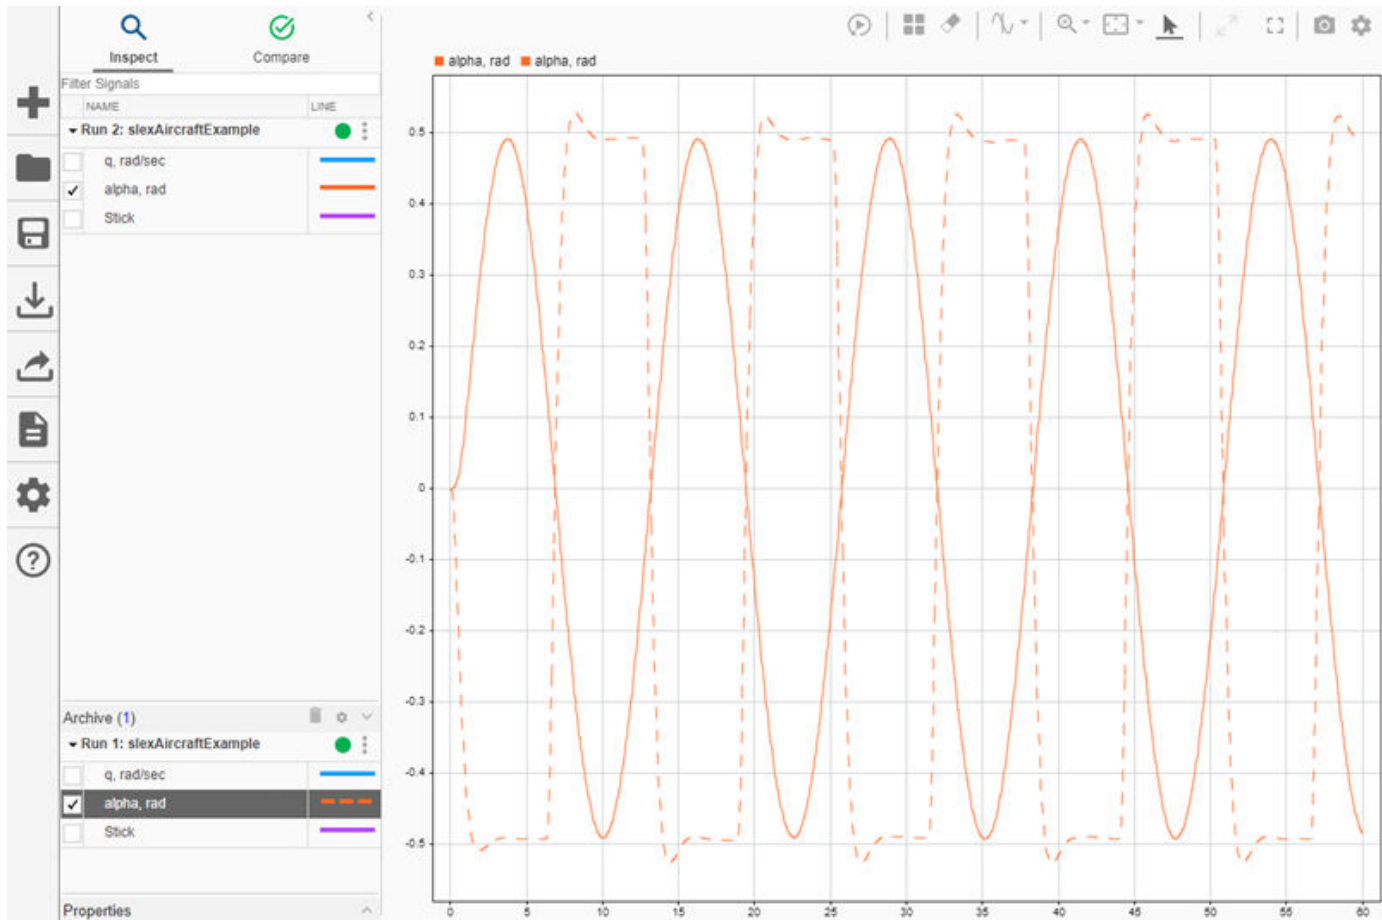
The Simulation Data Inspector manages incoming simulation data using the archive. By default, the previous run moves to the archive when you start a new simulation. You can plot signals from the archive, or you can drag runs of interest back into the work area.

Import Data from the Workspace or a File

You can import data from the base workspace or from a file to view on its own or alongside simulation data. The Simulation Data Inspector supports all built-in data types and many data formats for importing data from the workspace. In general, whatever the format, sample values must be paired with sample times. The Simulation Data Inspector allows up to 8000 channels per signal in a run created from imported workspace data.

You can also import data from these types of files:

- MAT file
- CSV file — Format data as shown in "Import Data from a CSV File into the Simulation Data Inspector" (Simulink).
- Microsoft® Excel® file — Format data as described in "Microsoft Excel Import, Export, and Logging Format" (Simulink).

- MDF file — MDF file import is supported for Linux® and Windows® operating systems. The MDF file must have a `.mdf`, `.mf4`, `.mf3`, `.data`, or `.dat` file extension and contain data with only integer and floating data types.
- ULG file — Flight log data import requires a UAV Toolbox license.

To import data from the workspace or from a file that is saved in a data or file format that the Simulation Data Inspector does not support, you can write your own workspace data or file reader to import the data using the `io.reader` class. You can also write a custom reader to use instead of the built-in reader for supported file types. For examples, see:

- "Import Data Using a Custom File Reader" (Simulink)
- "Import Workspace Variables Using a Custom Data Reader" (Simulink)

To import data, select the **Import** button in the Simulation Data Inspector.

In the Import dialog, you can choose to import data from the workspace or from a file. The table below the options shows data available for import. If you do not see your workspace variable or file contents in the table, that means the Simulation Data Inspector does not have a built-in or registered reader that supports that data. You can select which data to import using the check boxes, and you can choose whether to import that data into an existing run or a new run. To select all or none of the data, use the check box next to **NAME**.
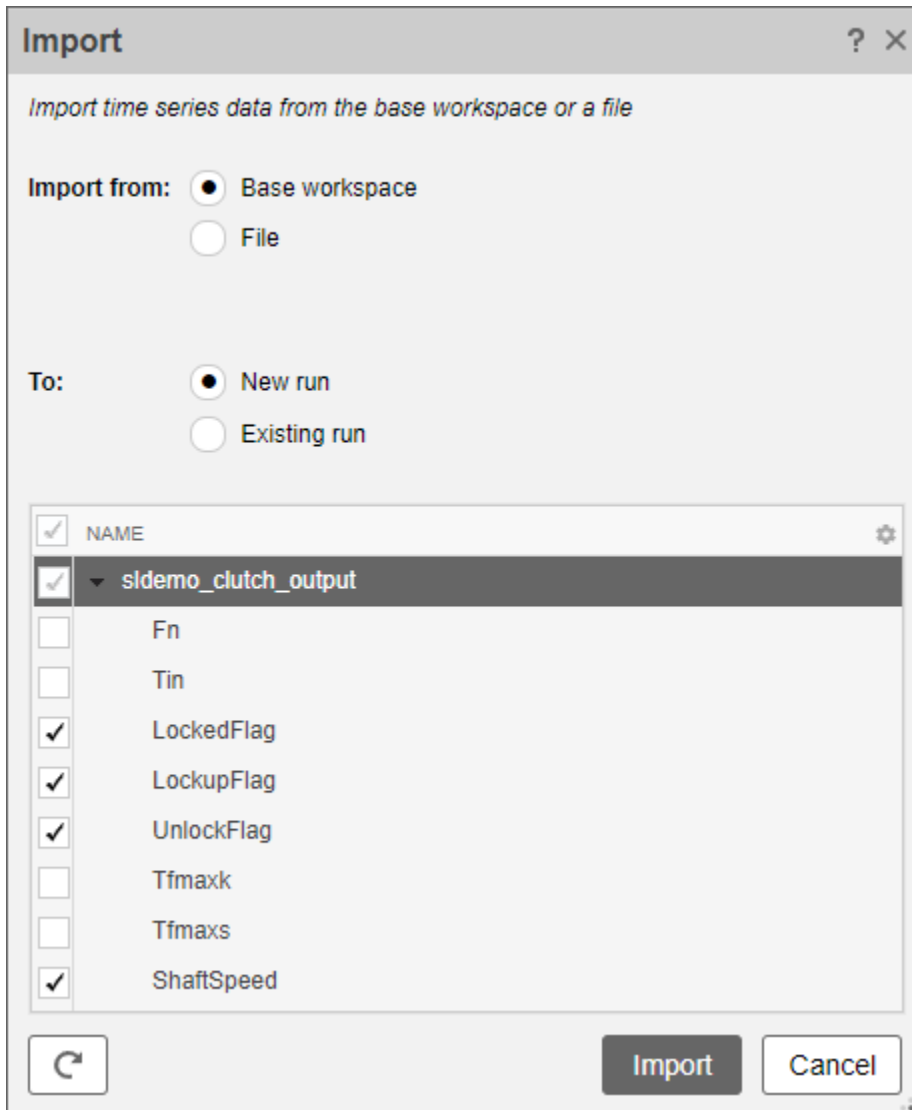
When you import data into a new run, the run always appears in the work area. You can manually move imported runs to the archive.

View Complex Data

To view complex data in the Simulation Data Inspector, import the data or log the signals to the Simulation Data Inspector. You can control how to visualize the complex signal using the **Properties** pane in the Simulation Data Inspector and in the **Instrumentation Properties** for the signal in the model. To access the **Instrumentation Properties** for a signal, right-click the logging badge for the signal and select **Properties**.

You can specify the **Complex Format** as Magnitude, Magnitude-Phase, Phase, or Real-Imaginary. If you select Magnitude-Phase or Real-Imaginary for the **Complex Format**, the Simulation Data Inspector plots both components of the signal when you select the check box for the signal. For signals in Real-Imaginary format, the **Line Color** specifies the color of the real component of the signal, and the imaginary component is a different shade of the **Line Color**. For example, the

Complex Signal displays the real component of the signal in light blue, matching the **Line Color** parameter, and the imaginary component is shown in a darker shade of blue.
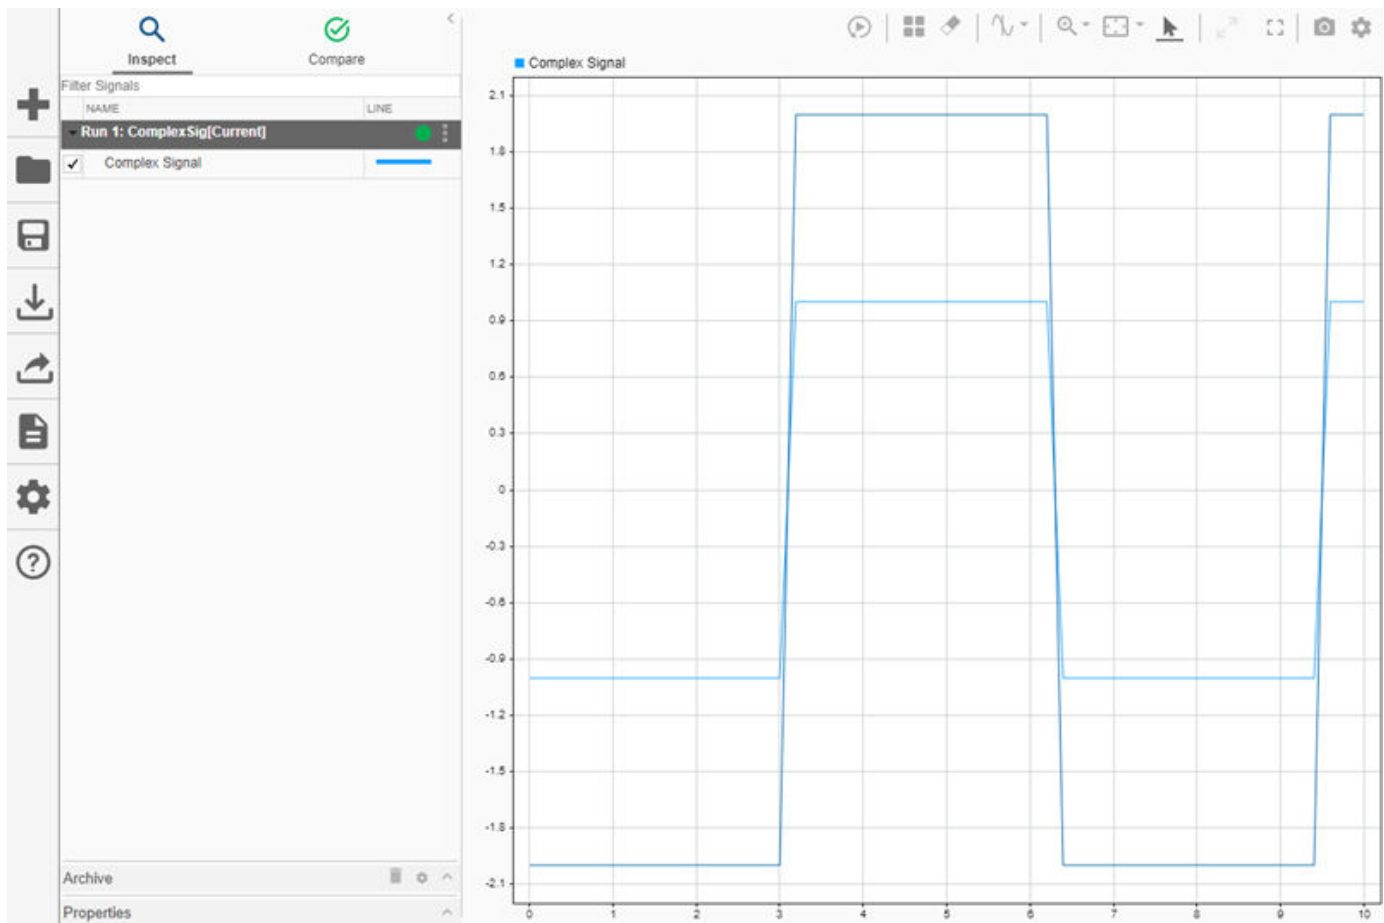
For signals in Magnitude-Phase format, the **Line Color** specifies the color of the magnitude component, and the phase is displayed in a different shade of the **Line Color**.

View String Data

You can log and view string data with your signal data in the Simulation Data Inspector. For example, consider this simple model. The value of the sine wave block controls whether the switch sends a string reading Positive or Negative to the output.
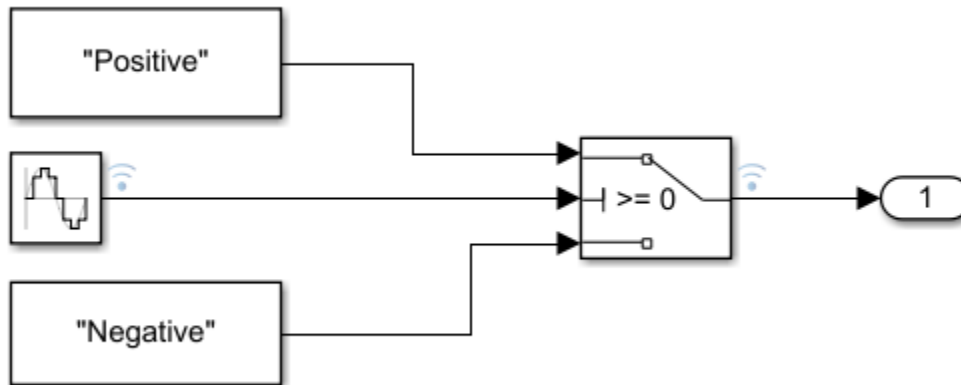
The plot shows the results of simulating the model. The string signal is shown at the bottom of the graphical viewing area. The value of the signal is displayed inside a band, and transitions in the string signal's value are marked with criss-crossed lines.
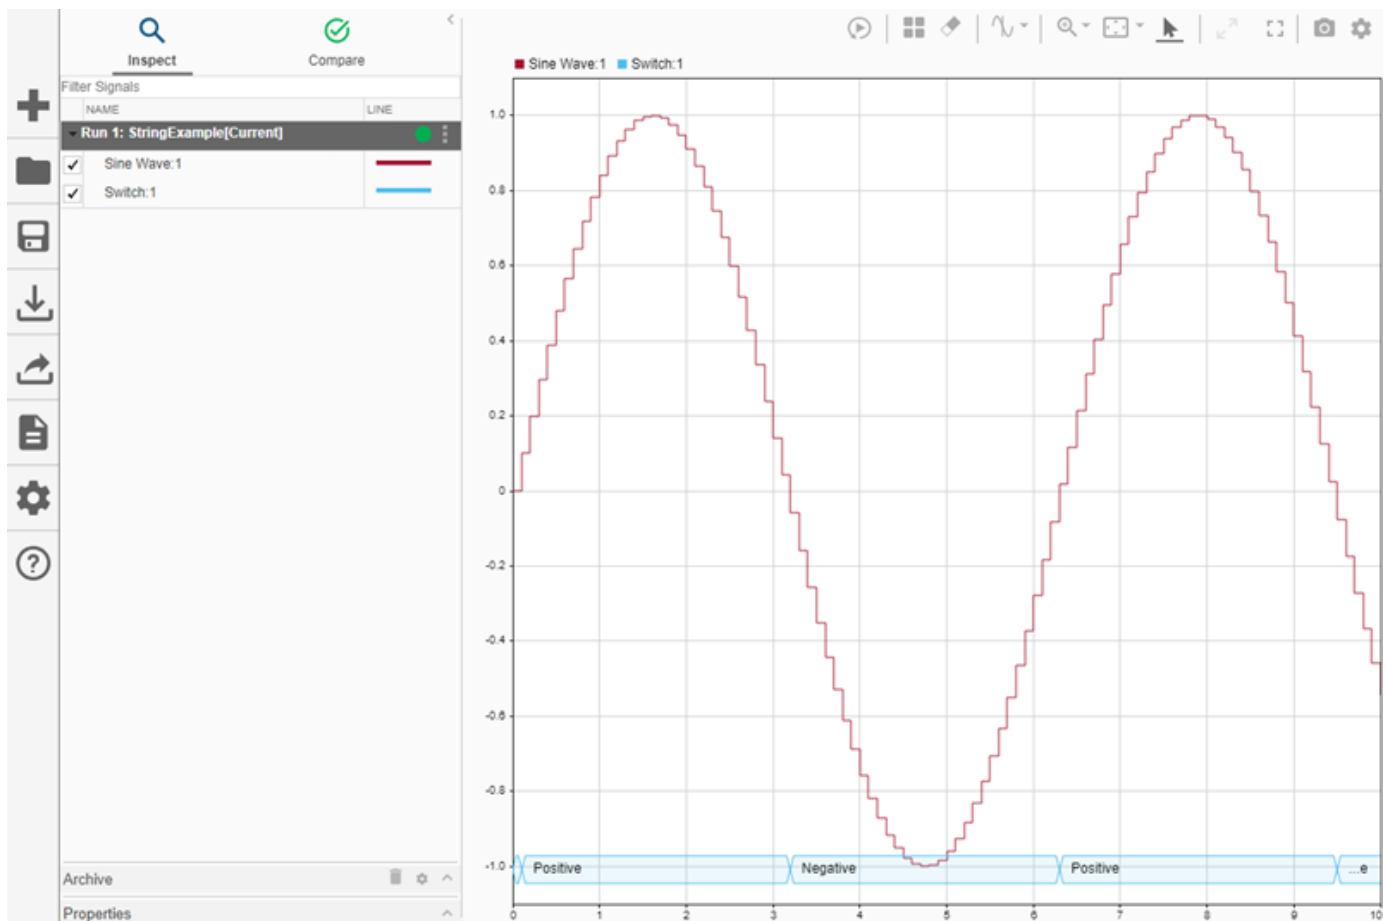
You can use cursors to inspect how the string signal values correspond with the sine signal's values.
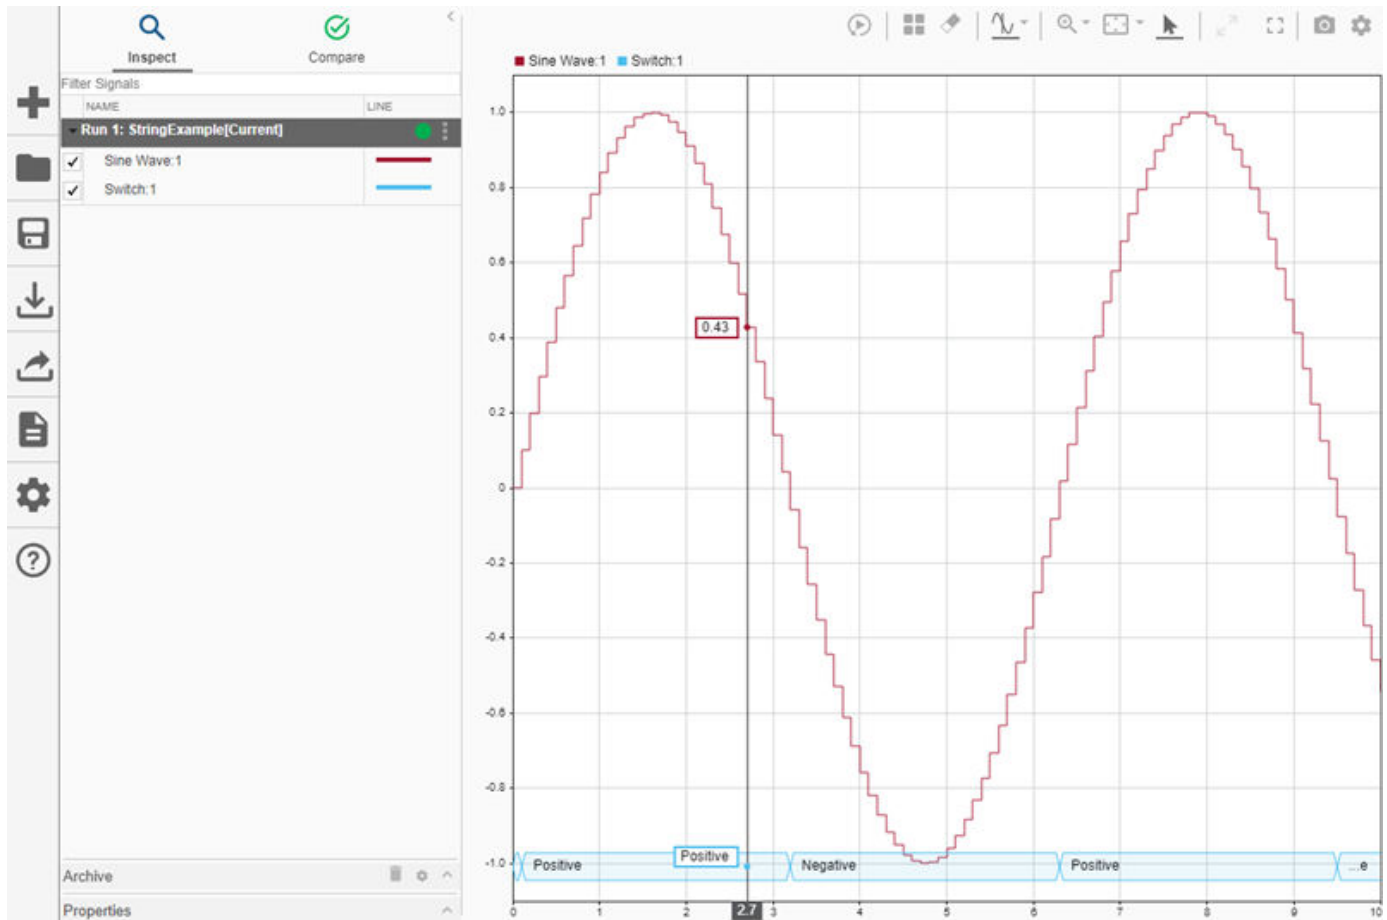
When you plot multiple string signals on a plot, the signals stack in the order they were simulated or imported, with the most recent signal positioned at the top. For example, you might consider the effect of changing the phase of the sine wave controlling the switch.
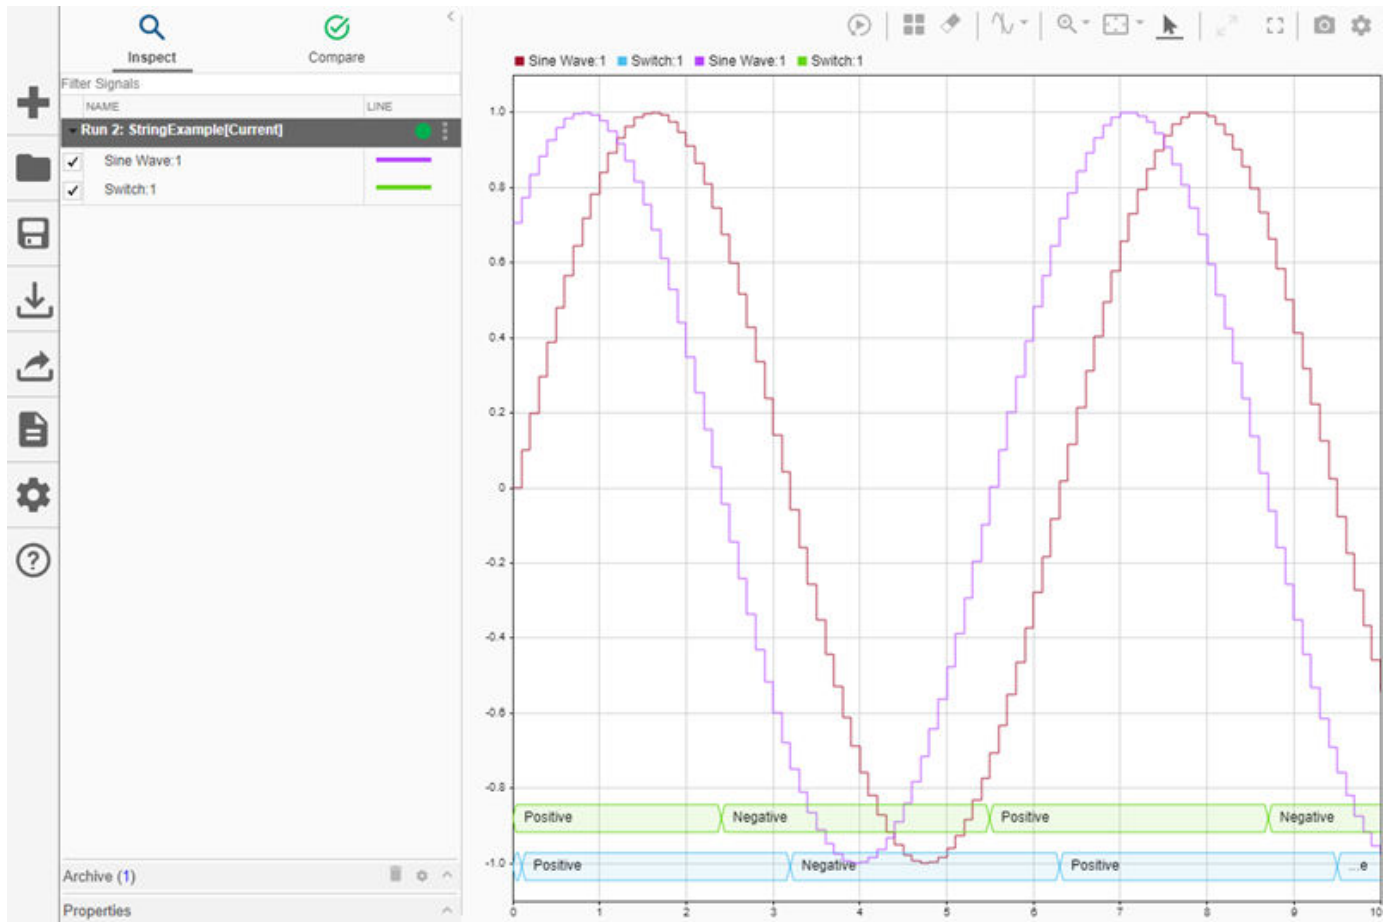
View Frame-Based Data

Processing data in frames rather than point by point provides a performance boost needed in some applications. To view frame-based data in the Simulation Data Inspector, you have to specify that the signal is frame-based in the **Instrumentation Properties** for the signal. To access the **Instrumentation Properties** dialog for a signal, right-click the signal's logging badge and select **Properties**. To specify a signal as frame-based, select **Columns as channels (frame based)** for **Input processing**.

View Event-Based Data

You can log or import event data to the Simulation Data Inspector. To view the logged event-based data, select the check box next to `Send: 1`. The Simulation Data Inspector displays the data as a stem plot, with each stem representing the number of events that occurred for a given sample time.
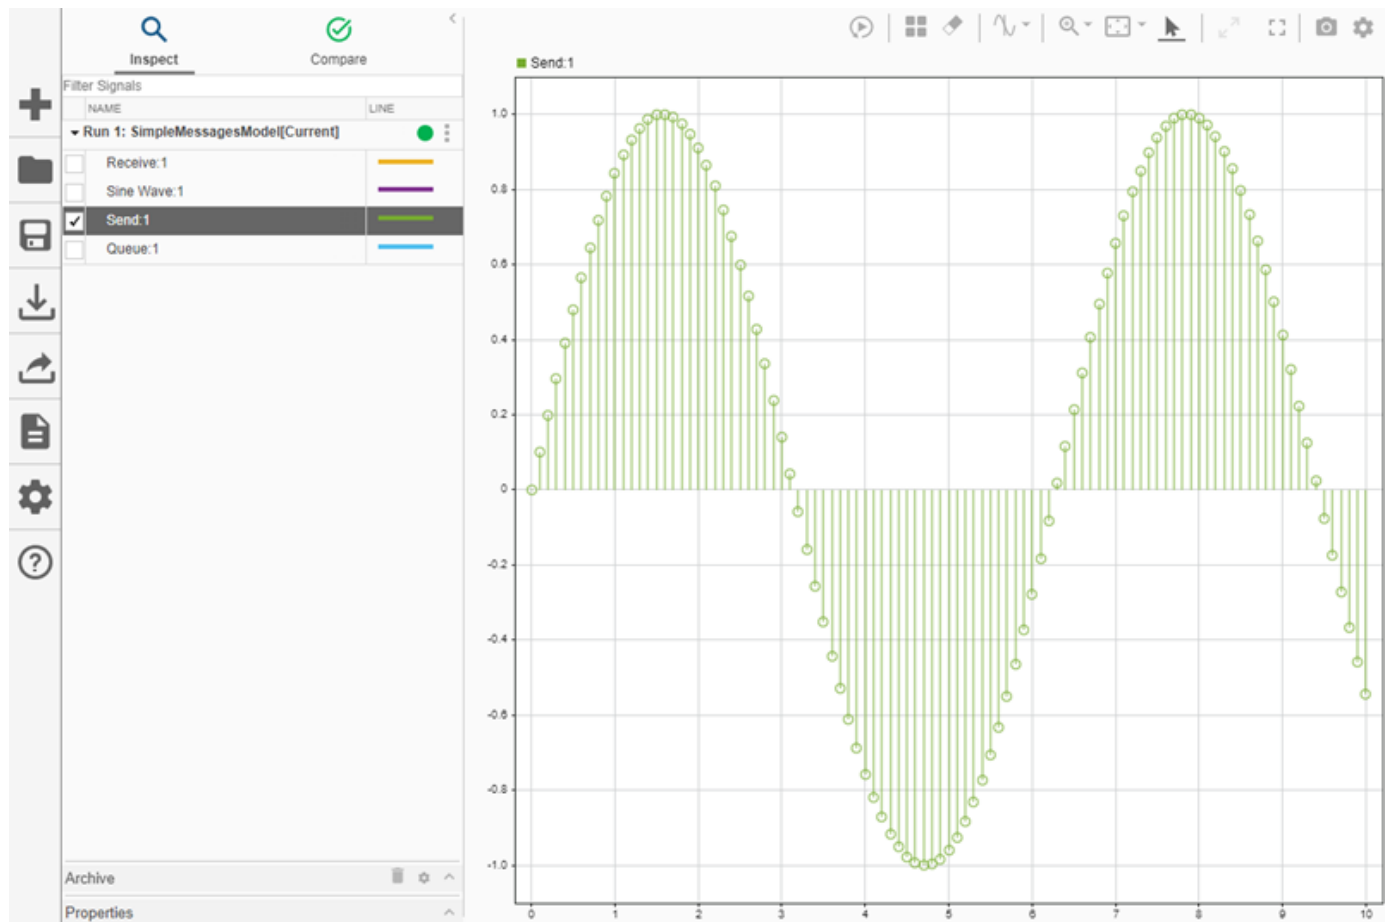
See Also

More About

- Inspect Simulation Data (Simulink)
- Compare Simulation Data (Simulink)
- Share Simulation Data Inspector Data and Views on page 5-36
- Decide How to Visualize Data (Simulink)
- Dataset Conversion for Logged Data (Simulink)

Import Data from a CSV File into the Simulation Data Inspector

To import data into the Simulation Data Inspector from a CSV file, format the data in the CSV file. Then, you can import the data using the Simulation Data Inspector UI or the `Simulink.sdi.createRun` function.
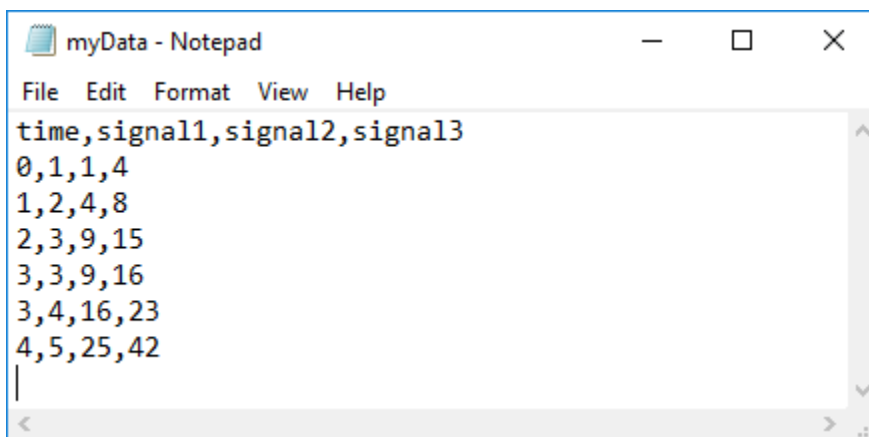
Tip When you want to import data from a CSV file where the data is formatted differently from the specification in this topic, you can write your own file reader for the Simulation Data Inspector using the `io.reader` class.

Basic File Format

In the simplest format, the first row in the CSV file is a header that lists the names of the signals in the file. The first column is time. The name for the time column must be `time`, and the time values must increase monotonically. The rows below the signal names list the signal values that correspond to each time step.

The import operation does not support time data that includes `Inf` or `NaN` values or signal data that includes `Inf` values. Empty or `NaN` signal values render as missing data. All built-in data types are supported.

Multiple Time Vectors

When your data includes signals with different time vectors, the file can include more than one time vector. Every time column must be named `time`. Time columns specify the sample times for signals to the right, up to the next time vector. For example, the first time column defines the time for `signal1` and `signal2`, and the second time column defines the time steps for `signal3`.

Signal columns must have the same number of data points as the associated time vector.

Signal Metadata

You can specify signal metadata in the CSV file to indicate the signal data type, units, interpolation method, block path, and port index. List metadata for each signal in rows between the signal name and the signal data. Label metadata according to this table.

Signal Property	Label	Value
Data type	`Type:`	Built-in data type.
Units	`Unit:`	Supported unit. For example, `Unit: m/s` specifies units of meters per second. For a list of supported units, enter `showunitslist` in the MATLAB Command Window.
Interpolation method	`Interp:`	`linear`, `zoh` for zero order hold, or `none`.
Block Path	`BlockPath:`	Path to the block that generated the signal.
Port Index	`PortIndex:`	Integer.

You can also import a signal with a data type defined by an enumeration class. Instead of using the `Type:` label, use the `Enum:` label and specify the value as the name of the enumeration class. The definition for the enumeration class must be saved on the MATLAB path.

When an imported file does not specify signal metadata, the Simulation Data Inspector assumes double data type and linear interpolation. You can specify the interpolation method as linear, zoh (zero-order hold), or none. If you do not specify units for the signals in your file, you can assign units to the signals in the Simulation Data Inspector after you import the file.

You can specify any combination of metadata for each signal. Leave a blank cell for signals with less specified metadata.

```
myData - Notepad                              —      □      ✕

File  Edit  Format  View  Help

time,signal1,signal2,time,signal3
 ,Interp: zoh, , ,Interp: zoh
 ,Type: int8,Type: int32
 ,Unit: m, , ,Unit: m/s
0,1,1,0,4
1,2,4,2,8
2,3,9,3,15
3,3,9,5,16
3,4,16
4,5,25
|
```

Import Data from a CSV File

You can import data from a CSV file using the Simulation Data Inspector UI or using the `Simulink.sdi.createRun` function.

To import data using the UI, open the Simulation Data Inspector using the `Simulink.sdi.view` function or the **Data Inspector** button in the Simulink™ toolstrip. Then, click **Import** .

In the Import dialog, select the option to import data from a file and navigate in the file system to select the file. After you select the file, data available for import shows in the table. You can choose which signals to import and whether to import them to a new or existing run. This example imports all available signals to a new run. To select all or none of the signals, select or clear the check box next to NAME. After selecting the options, click the **Import** button.
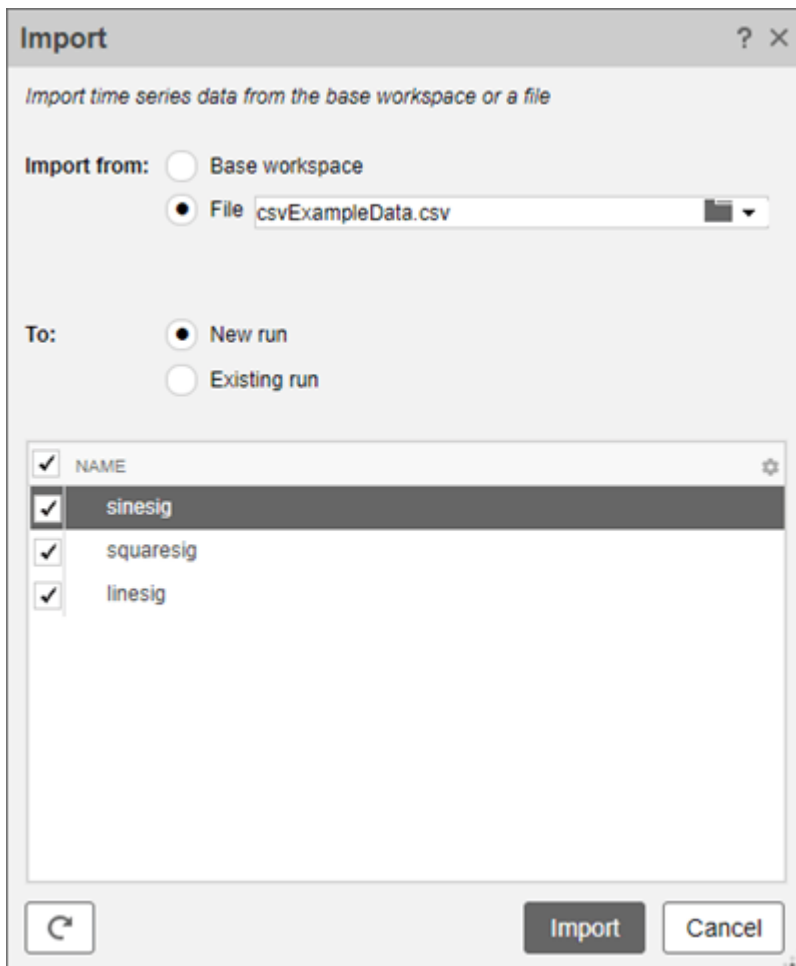
When you import data into a new run using the UI, the new run name includes the run number followed by `Imported_Data`.

When you import data programmatically, you can specify the name of the imported run.

```
csvRunID = Simulink.sdi.createRun('CSV File Run','file','csvExampleData.csv');
```

See Also

Functions
`Simulink.sdi.createRun`

More About

- "View Data in the Simulation Data Inspector" (Simulink)
- "Microsoft Excel Import, Export, and Logging Format" (Simulink)
- "Import Data Using a Custom File Reader" (Simulink)

Microsoft Excel Import, Export, and Logging Format

Using the Simulation Data Inspector or Simulink Test, you can import data from a Microsoft Excel file or export data to a Microsoft Excel file. You can also log data to an Excel file using the Record block. The Simulation Data Inspector, Simulink Test, and the Record block all use the same file format, so you can use the same Microsoft Excel file with multiple applications.

Tip When the format of the data in your Excel file does not match the specification in this topic, you can write your own file reader to import the data using the `io.reader` class.

Basic File Format

In the simplest format, the first row in the Excel file is a header that lists the names of the signals in the file. The first column is time. The name for the time column must be `time`, and the time values must increase monotonically. The rows below the signal names list the signal values that correspond to each time step.

	A	B	C	D
1	time	signal1	signal2	signal3
2	0	1	1	4
3	1	2	4	8
4	2	3	9	15
5	3	3	9	16
6	3	4	16	23
7	4	5	25	42

The import operation does not support time data that includes `Inf` or `NaN` values or signal data that includes `Inf` values. Empty or `NaN` signal values imported from the Excel file render as missing data in the Simulation Data Inspector. All built-in data types are supported.

Multiple Time Vectors

When your data includes signals with different time vectors, the file can include more than one time vector. Every time column must be named `time`. Time columns specify the sample times for signals to the right, up to the next time vector. For example, the first time column defines the time for `signal1` and `signal2`, and the second time column defines the time steps for `signal3`.

	A	B	C	D	E
1	time	signal1	signal2	time	signal3
2	0	1	1	0	4
3	1	2	4	2	8
4	2	3	9	3	15
5	3	3	9	5	16
6	3	4	16		
7	4	5	25		

Signal columns must have the same number of data points as the associated time vector.

Signal Metadata

The file can include metadata for signals such as data type, units, and interpolation method. The metadata is used to determine how to plot the data, how to apply unit and data conversions, and how to compute comparison results. For more information about how metadata is used in comparisons, see "How the Simulation Data Inspector Compares Data" (Simulink).

Metadata for each signal is listed in rows between the signal names and the signal data. You can specify any combination of metadata for each signal. Leave a blank cell for signals with less specified metadata.

	A	B	C	D	E
1	time	signal1	signal2	time	signal3
2		Interp: zoh			Interp: zoh
3		Type: int8	Type: int32		
4		Unit: m			Unit: m/s
5	0	1	1	0	4
6	1	2	4	2	8
7	2	3	9	3	15
8	3	3	9	5	16
9	3	4	16		
10	4	5	25		

Label each piece of metadata according to this table. The table also indicates which tools and operations support each piece of metadata. When an imported file does not specify signal metadata, `double` data type, `linear` interpolation, and `union` synchronization are used.

Property Descriptions

Signal Property	Label	Values	Simulation Data Inspector Import	Record Block Logging and Simulation Data Inspector Export	Simulink Test Import and Export
Data type	`Type:`	Built-in data type.	**Supported**	**Supported**	**Supported**
Units	`Unit:`	Supported unit. For example, `Unit: m/s` specifies units of meters per second. For a list of supported units, enter `showunitslist` in the MATLAB Command Window.	**Supported**	**Supported**	**Supported**
Interpolation method	`Interp:`	`linear`, `zoh` for zero order hold, or `none`.	**Supported**	**Supported**	**Supported**
Synchronization method	`Sync:`	`union` or `intersection`.	**Supported**	**Not Supported** *Metadata not included in exported file.*	**Supported**
Relative tolerance	`RelTol:`	Percentage, represented as a decimal. For example, `RelTol: 0.1` specifies a 10% relative tolerance.	**Supported**	**Not Supported** *Metadata not included in exported file.*	**Supported**
Absolute tolerance	`AbsTol:`	Numeric value.	**Supported**	**Not Supported** *Metadata not included in exported file.*	**Supported**
Time tolerance	`TimeTol:`	Numeric value, in seconds.	**Supported**	**Not Supported** *Metadata not included in exported file.*	**Supported**

Signal Property	Label	Values	Simulation Data Inspector Import	Record Block Logging and Simulation Data Inspector Export	Simulink Test Import and Export
Leading tolerance	`LeadingTol:`	Numeric value, in seconds.	**Supported** *Only visible in Simulink Test.*	**Not Supported** *Metadata not included in exported file.*	**Supported**
Lagging tolerance	`LaggingTol:`	Numeric Value, in seconds.	**Supported** *Only visible in Simulink Test.*	**Not Supported** *Metadata not included in exported file.*	**Supported**
Block Path	`BlockPath:`	Path to the block that generated the signal.	**Supported**	**Supported**	**Supported**
Port Index	`PortIndex:`	Integer.	**Supported**	**Supported**	**Supported**
Name	`Name:`	Signal name	**Supported**	**Not Supported** *Metadata not included in exported file.*	**Supported**

User-Defined Data Types

In addition to built-in data types, you can use other labels in place of the `DataType:` label to specify fixed-point, enumerated, alias, and bus data types.

Property Descriptions

Data Type	Label	Values	Simulation Data Inspector Import	Record Block Logging and Simulation Data Inspector Export	Simulink Test Import and Export
Enumeration	`Enum:`	Name of the enumeration class.	**Supported** *Enumeration class definition must be saved on the MATLAB path.*	**Supported** *Enumeration class definition must be saved on the MATLAB path.*	**Supported** *Enumeration class definition must be saved on the MATLAB path.*
Alias	`Alias:`	Name of a `Simulink.AliasType` object in the MATLAB workspace.	**Supported** *For matrix and complex signals, specify the alias data type on the first channel.*	**Not Supported**	**Supported** *For matrix and complex signals, specify the alias data type on the first channel.*
Fixed-point	`Fixdt:`	• `fixdt` constructor. • Name of a `Simulink. NumericType` object in the MATLAB workspace. • Name of a fixed-point data type as described in "Fixed-Point Numbers in Simulink" (Fixed-Point Designer).	**Supported**	**Not Supported**	**Supported**
Bus	`Bus:`	Name of a `Simulink.Bus` object in the MATLAB workspace.	**Supported**	**Not Supported**	**Supported**

When you specify the type using the name of a `Simulink.Bus` object and the object is not in the MATLAB workspace, the data still imports from the file. However, individual signals in the bus use data types described in the file rather than data types defined in the `Simulink.Bus` object.

Complex, Multidimensional, and Bus Signals

You can import and export complex, multidimensional, and bus signals using an Excel file. The signal name for a column of data indicates whether that data is part of a complex, multidimensional, or bus signal. Excel file import and export do not support array of bus signals.

Note When you export data from a nonvirtual bus with variable-size signals to an Excel file, the variable-size signal data is expanded to individual channels, and the hierarchical nature of the data is lost. Data imported from this file is returned as a flat list.

Multidimensional signal names include index information in parentheses. For example, the signal name for a column might be `signal1(2,3)`. When you import data from a file that includes multidimensional signal data, elements in the data not included in the file take zero sample values with the same data type and complexity as the other elements.

Complex signal data is always in real-imaginary format. Signal names for columns containing complex signal data include `(real)` and `(imag)` to indicate which data each column contains. When you import data from a file that includes imaginary signal data without specifying values for the real component of that signal, the signal values for the real component default to zero.

Multidimensional signals can contain complex data. The signal name includes the indication for the index within the multidimensional signal and the real or imaginary tag. For example, `signal1(1,3)`
`(real)`.

Dots in signal names specify the hierarchy for bus signals. For example:

- `bus.y.a`
- `bus.y.b`
- `bus.x`

	A	B	C	D	E
1	time	bus.y.a	bus.y.b	time	bus.x
2		Interp: zoh			Interp: zoh
3		Type: int8	Type: int32		
4		Unit: m			Unit: m/s
5	0	1	1	0	4
6	1	2	4	2	8
7	2	3	9	3	15
8	3	3	9	5	16
9	3	4	16		
10	4	5	25		

Tip When the name of your signal includes characters that could make it appear as though it were part of a matrix, complex signal, or bus, use the `Name` metadata option to specify the name you want the imported signal to use in the Simulation Data Inspector and Simulink Test.

Function-Call Signals

Signal data specified in columns before the first time column is imported as one or more function-call signals. The data in the column specifies the times at which the function-call signal was enabled. The imported signals have a value of 1 for the times specified in the column. The time values for function-call signals must be double, scalar, and real, and must increase monotonically.

When you export data from the Simulation Data Inspector, function-call signals are formatted the same as other signals, with a time column and a column for signal values.

Simulation Parameters

You can import data for parameter values used in simulation. In the Simulation Data Inspector, the parameter values are shown as signals. Simulink Test uses imported parameter values to specify values for those parameters in the tests it runs based on imported data.

Parameter data is specified using two or three columns. The first column specifies the parameter names, with the cell in the header row for that column labeled `Parameter:`. The second column specifies the value used for each parameter, with the cell in the header row labeled `Value:`. Parameter data may also include a third column that contains the block path associated with each parameter, with the cell in the header row labeled `BlockPath:`. Specify names, values, and block paths for parameters starting in the first row that contains signal data, below rows used to specify signal metadata. For example, this file specifies values for two parameters, X and Y.

	A	B	C	D	E	F	G
1	time	signal1	signal2	time	signal3	Parameter:	Value:
2		Interp: zoh			Interp: zoh		
3		Type: int8	Type: int32				
4		Unit: m			Unit: m/s		
5	0	1	1	0	4	X	2
6	1	2	4	2	8	Y	1.2
7	2	3	9	3	15		
8	3	3	9	5	16		
9	3	4	16				
10	4	5	25				

Multiple Runs

You can include data for multiple runs in a single file. Within a sheet, you can divide data into runs by labeling data with a simulation number and a source type, such as `Input` or `Output`. Specify the simulation number and source type as additional signal metadata, using the label `Simulation:` for the simulation number and the label `Source:` for the source type. The Simulation Data Inspector uses the simulation number and source type only to determine which signals belong in each run. Simulink Test uses the information to define inputs, parameters, and acceptance criteria for tests to run based on imported data.

You do not need to specify the simulation number and output type for every signal. Signals to the right of a signal with a simulation number and source use the same simulation number and source

until the next signal with a different source or simulation number. For example, this file defines data for two simulations and imports into four runs in the Simulation Data Inspector:

- **Run 1** contains `signal1` and `signal2`.
- **Run 2** contains `signal3`, X, and Y.
- **Run 3** contains `signal4`.
- **Run 4** contains `signal5`.

	A	B	C	D	E	F	G	H	I	J
1	time	signal1	signal2	time	signal3	Parameter:	Values:	time	signal4	signal5
2		Interp: zoh			Interp: zoh					
3		Type: int8	Type: int32							
4		Unit: m			Unit: m/s					
5		Simulation: 1							Simulation: 2	
6		Source: Input			Source: Output				Source: Input	Source: Output
7	0	1	1	0	4	X	2	1	2	1
8	1	2	4	2	8	Y	1.2	2	6	3
9	2	3	9	3	15			3	4	5
10	3	3	9	5	16			4	8	7
11	3	4	16					5	10	2
12	4	5	25							

You can also use sheets within the Microsoft Excel file to divide the data into runs and tests. When you do not specify simulation number and source information, the data on each sheet is imported into a separate run in the Simulation Data Inspector. When you export multiple runs from the Simulation Data Inspector, the data for each run is saved on a separate sheet. When you import a Microsoft Excel file that contains data on multiple sheets into Simulink Test, you are prompted to specify how to import the data.

See Also
`Simulink.sdi.createRun` | `Simulink.sdi.exportRun`

More About
- "View Data in the Simulation Data Inspector" (Simulink)
- "Import Data from a CSV File into the Simulation Data Inspector" (Simulink)
- "Import Data Using a Custom File Reader" (Simulink)

Configure the Simulation Data Inspector

The Simulation Data Inspector supports a wide range of use cases for analyzing and visualizing data. You can modify preferences in the Simulation Data Inspector to match your visualization and analysis requirements. The preferences that you specify persist between MATLAB sessions.

By specifying preferences in the Simulation Data Inspector, you can configure options such as:

- How signals and metadata are displayed.
- Which data automatically imports from parallel simulations.
- Where prior run data is retained and how much prior data to store.
- How much memory is used during save operations.
- The system of units used to display signals.

To open the Simulation Data Inspector preferences, click Preferences.

Note You can restore all preferences in the Simulation Data Inspector to default values by clicking **Restore Defaults** in the Preferences menu or by using the `Simulink.sdi.clearPreferences` function.

Logged Data Size and Location

By default, simulation data logs to disk with data loaded into memory on demand, and the maximum size of logged data is constrained only by available disk space. You can use the **Disk Management** settings in the Simulation Data Inspector to directly control the size and location of logged data.

The **Record mode** setting specifies whether logged data is retained after simulation. When you change the **Record mode** setting to **View during simulation only**, no logged data is available in the Simulation Data Inspector or the workspace after the simulation completes. Only use this mode when you do not want to save logged data. The **Record mode** setting reverts to **View and record data** each time you start MATLAB. Changing the **Record mode** setting can affect other applications, such as visualization tools. For details, see "View Data Only During Simulation" (Simulink).

To directly limit the size of logged data, you can specify a minimum amount of free disk space or a maximum size for the logged data. By default, logged data must leave at least 100 MB of free disk space with no maximum size limit. Specify the required disk space and maximum size in GB, and specify 0 to apply no disk space requirement or no maximum size limit.

When you specify a minimum disk space requirement or a maximum size for logged data, you can also specify whether to prioritize retaining data from the current simulation or data from prior simulations when approaching the limit. By default, the Simulation Data Inspector prioritizes retaining data for the current run by deleting data for prior runs. To prioritize retaining prior data, change the **When low on disk space** setting to **Keep prior runs and stop recording**. You see a warning message when prior runs are deleted and when recording is disabled. If recording is disabled due to the size of logged data, you need to change the **Record Mode** back to **View and**

record data to continue logging data, after you have freed up disk space. For more information, see "Specify a Minimum Disk Space Requirement or Maximum Size for Logged Data" (Simulink).

The **Storage Mode** setting specifies whether to log data to disk or to memory. By default, data logs to disk. When you configure a parallel worker to log data to memory, data transfer back to the host is not supported. Logging data to memory is not supported for rapid accelerator simulations or models deployed using Simulink Compiler™.

You can also specify the location of the temporary file that stores logged data. By default, data logs to the temporary files directory on your computer. You may change the file location when you need to log large amounts of data and a secondary drive provides more storage capacity. Logging data to a network location can degrade performance.

Programmatic Use

You can programmatically configure and check each preference value.

Preference	Functions
Record mode	Simulink.sdi.setRecordData Simulink.sdi.getRecordData
Required Free Space	Simulink.sdi.setRequiredFreeSpace Simulink.sdi.getRequiredFreeSpace
Max Disk Usage	Simulink.sdi.setMaxDiskUsage Simulink.sdi.getMaxDiskUsage
When low on disk space	Simulink.sdi.setDeleteRunsOnLowSpace Simulink.sdi.getDeleteRunsOnLowSpace
Storage Mode	Simulink.sdi.setStorageMode Simulink.sdi.getStorageMode
Storage Location	Simulink.sdi.setStorageLocation Simulink.sdi.getStorageLocation

Archive Behavior and Run Limit

When you run multiple simulations in a single MATLAB session, the Simulation Data Inspector retains results from each simulation so you can analyze the results together. Use the Simulation Data Inspector archive to manage runs in the user interface and control the number of runs the Simulation Data Inspector retains.

You can configure a limit for the number of runs to retain in the archive and whether the Simulation Data Inspector automatically moves prior runs into the archive.

Manage Runs Using the Archive

By default, the Simulation Data Inspector automatically archives simulation runs. When you simulate a model, the prior simulation run moves to the archive, and the Simulation Data Inspector updates the view to show data for aligned signals in the current run.

The archive does not impose functional limitations on the runs and signals it contains. You can plot signals from the archive, and you can use runs and signals in the archive in comparisons. You can drag runs of interest from the archive to the work area and vice versa whether **Automatically Archive** is selected or disabled.

To prevent the Simulation Data Inspector from automatically moving prior simulations runs to the archive, clear the **Automatically archive** setting. With automatic archiving disabled, the Simulation Data Inspector does not move prior runs into the **Archive** pane or automatically update plots to display data from the current simulation.

Tip To manually delete the contents of the archive, click Delete archived runs 🗑.

Control Number of Runs Retained in Simulation Data Inspector

You can specify a limit for the number of runs to retain in the archive. When the number of runs in the archive reaches the limit, the Simulation Data Inspector deletes runs in the archive on a first-in, first-out basis.

The run limit applies only to runs in the archive. For the Simulation Data Inspector to automatically limit the data it retains by deleting old runs, select **Automatically archive** and specify a size limit.

By default, the Simulation Data Inspector retains the last 20 runs moved to the archive. To remove the limit, select **No limit**. To specify the maximum number of runs to store in the archive, select **Last n runs** and enter the limit. A warning occurs if you specify a limit that would delete runs already in the archive.

Programmatic Use

You can programmatically configure and check the archive behavior and run limit.

Preference	Functions

Automatically archive	`Simulink.sdi.setAutoArchiveMode`
	
	`Simulink.sdi.getAutoArchiveMode`

Size	`Simulink.sdi.setArchiveRunLimit`
	
	`Simulink.sdi.getArchiveRunLimit`

Incoming Run Names and Location

You can configure how the Simulation Data Inspector handles incoming runs from import or simulation. You can choose whether new runs are added at the top of the work area or the bottom and specify a naming rule to use for runs created from simulation.

By default, the Simulation Data Inspector adds new runs below prior runs in the work area. The **Archive** settings also affect the location of runs. By default, prior runs are moved to the archive when a new simulation run is created.

The run naming rule is used to name runs created from simulation. You can create the run naming rule using a mix of literal text that is used in the run name as-is and one or more tokens that represent metadata about the run. By default, the Simulation Data Inspector names runs using the run index and model name: Run <run_index>: <model_name>.

Tip To rename an existing run, double-click the name in the work area and enter the new name, or modify the run name in the **Properties** pane.

Programmatic Use

You can programmatically configure and check incoming run names and locations.

Preference	Functions
Add New Runs	Simulink.sdi.setAppendRunToTop
	Simulink.sdi.getAppendRunToTop
Naming Rule	Simulink.sdi.setRunNamingRule
	Simulink.sdi.getRunNamingRule
	Simulink.sdi.resetRunNamingRule

Signal Metadata to Display

You can control which signal metadata is displayed in the work area of the **Inspect** pane and in the results section on the **Compare** pane in the Simulation Data Inspector. You specify the metadata to display separately for each pane using the **Table Columns** preferences in the **Inspect** and **Compare** sections of the Preferences dialog, respectively.

Inspect Pane

By default, the signal name and the line style and color used to plot the signal are displayed on the **Inspect** pane. To display different or additional metadata in the work area on the **Inspect** pane, select the check box next to each piece of metadata you want to display in the **Table Columns** preference in the **Inspect** section. You can always view complete metadata for the selected signal in the **Inspect** pane using the **Properties** pane.

Note Metadata displayed in the work area on **Inspect** pane is included when you generate a report of plotted signals. You can also specify metadata to include in the report regardless of what is displayed in the work area when you create the report programmatically using the Simulink.sdi.report function.

Compare Pane

By default, the **Compare** pane shows the signal name, the absolute and relative tolerances used in the signal comparison, and the maximum difference from the comparison result. To display different or additional metadata in the results on the **Compare** pane, select the check box next to each piece of metadata you want to display in the **Table Columns** preference in the **Compare** section. You can always view complete metadata for the signals compared for a selected signal result using the **Properties** pane, where metadata that differs between the compared signals is highlighted. Signal metadata displayed on the **Compare** pane does not affect the contents of comparison reports.

Signal Selection on the Inspect Pane

You can configure how you select signals to plot on the selected subplot in the Simulation Data Inspector. By default, you use check boxes next to each signal to plot. You can also choose to plot signals based on selection in the work area. Use **Check Mode** when creating views and visualizations that represent findings and analysis of a data set. Use **Browse Mode** to quickly view and analyze data sets with a large number of signals.

For more information about creating visualizations using **Check Mode**, see "Create Plots Using the Simulation Data Inspector" (Simulink).

For more information about using **Browse Mode**, see "Visualize Many Logged Signals" (Simulink).

Note To use **Browse Mode**, your layout must include only **Time Plot** visualizations.

How Signals Are Aligned for Comparison

When you compare runs using the Simulation Data Inspector, the comparison algorithm pairs signals for signal comparison through a process called alignment. You can align signals between the compared runs using one or more of the signal properties shown in the table.

Property	Description
Data Source	Path of the variable in the MATLAB workspace for data imported from the workspace
Path	Block path for the source of the data in its model
SID	Automatically assigned Simulink identifier
Signal Name	Name of the signal

You can specify the priority for each piece of metadata used for alignment. The **Align By** field specifies the highest priority property used to align signals. The priority drops with each subsequent **Then By** field. You must specify a primary alignment property in the **Align By** field, but you can leave any number of **Then By** fields blank.

By default, the Simulation Data Inspector aligns signals between runs according to this flow chart.
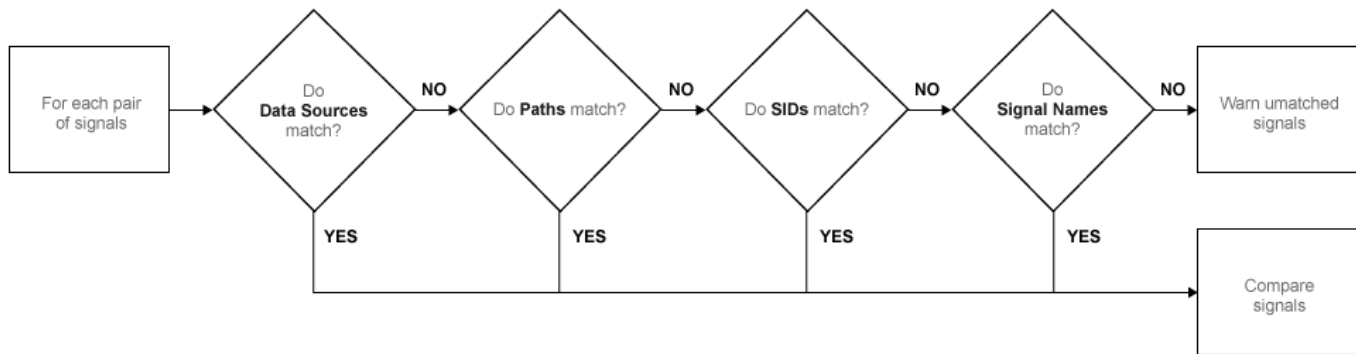
For more information about configuring comparisons in the Simulation Data Inspector, see "How the Simulation Data Inspector Compares Data" (Simulink).

Colors Used to Display Comparison Results

You can configure the colors used to display comparison results using the Simulation Data Inspector preferences. You can specify whether to use the signal color from the **Inspect** pane or a fixed color for the baseline and compared signals. You can also choose colors for the tolerance and the difference signal.

By default, the Simulation Data Inspector displays comparison results using fixed colors for the baseline and compared signals. Using a fixed color allows you to avoid the baseline signal color and compared signal color being either the same or too similar to distinguish.

Signal Grouping

You can specify how to group signals within a run in the Simulation Data Inspector. The preferences apply to both the **Inspect** and **Compare** panes and comparison reports. You can group signals by:

*   Domain — Signal type. For example, signals created by signal logging have a domain of `Signal`, while signals created from logging model outputs have a domain of `Outports`.
*   Physical System Hierarchy — Signal Simscape™ physical system hierarchy. The option to group by physical system hierarchy is available when you have a Simscape license.
*   Data Hierarchy — Signal location within structured data. For example, data hierarchy grouping reflects the hierarchy of a bus.
*   Model Hierarchy — Signal location within model hierarchy. Grouping by model hierarchy can be helpful when you log data from a model that includes model or subsystem references.

Grouping signals adds rows for the hierarchical nodes, which you can expand to show the signals within that node. By default, the Simulation Data Inspector groups signals by domain, then by physical system hierarchy (if you have a Simscape license), and then by data hierarchy.

To remove grouping and display a flat list of signals in each run, select `None` for all grouping options.

Programmatic Use

To specify how to group signals programmatically, use the `Simulink.sdi.setTableGrouping` function.

Data to Stream from Parallel Simulations

When you run parallel simulations using the `parsim` function, you can stream logged simulation data to the Simulation Data Inspector. A dot next to the run name in the **Inspect** pane indicates the status of the simulation that corresponds to the run, so you can monitor simulation progress while visualizing the streamed data. You can control whether data streams from a parallel simulation based on the type of worker the data comes from.

By default, the Simulation Data Inspector is configured for manual import of data from parallel workers. You can use the Simulation Data Inspector programmatic interface to inspect the data on the worker and decide whether to send it to the client Simulation Data Inspector for further analysis. To manually move data from a parallel worker to the Simulation Data Inspector, use the `Simulink.sdi.sendWorkerRunToClient` function.

You may want to automatically stream data from parallel simulations that run on local workers or on local and remote workers. Streaming data from both local and remote workers may affect simulation performance, depending on how many simulations you run and how much data you log. When you choose to stream data from local workers or all parallel workers, all logged simulation data automatically shows in the Simulation Data Inspector.

Programmatic Use

You can configure Simulation Data Inspector support for parallel worker data programmatically using the `Simulink.sdi.enablePCTSupport` function.

Options for Saving and Loading Session Files

You can specify a maximum amount of memory to use while loading or saving a session file. By default, the Simulation Data Inspector uses a maximum of 100 MB of memory when you load or save a session file. You can specify a memory use limit as low as 50 MB.

To reduce the size of the saved session file, you can specify a compression option.

- `None` — Do not compress saved data.
- `Normal` — Compress the saved file as much as possible.
- `Fastest` — Compress the saved file less than `Normal` compression for faster save time.

Signal Display Units

Signals in the Simulation Data Inspector have two units properties: stored units and display units. The stored units represent the units of the data saved to disk. The display units specify how the Simulation Data Inspector displays the data. You can configure the Simulation Data Inspector to use a system of units to define the display units for all signals. You can choose either the **SI** or **US Customary** system of units, or you can display data using its stored units.

When you use a system of units to define display units for signals in the Simulation Data Inspector, the display units update for any signal with display units that are not valid for that unit system. For example, if you select **SI** units, the display units for a signal may update from `ft` to `m`.

Note The system of units you choose to use in the Simulation Data Inspector does not affect the stored units for any signal. You can convert the stored units for a signal using the `convertUnits` function. Conversion may result in loss of precision.

In addition to selecting a system of units, you can specify override units so that all signals of a given measurement type are displayed using consistent units. For example, if you want to visualize all signals that represent weight using units of `kg`, specify `kg` as an override unit.

Tip For a list of units supported by Simulink, enter `showunitslist` in the MATLAB Command Window.

You can also modify the display units for a specific signal using the **Properties** pane. For more information, see "Modify Signal Properties in the Simulation Data Inspector" (Simulink).

Programmatic Use

Configure the unit system and override units using the `Simulink.sdi.setUnitSystem` function. You can check the current units preferences using the `Simulink.sdi.getUnitSystem` function.

See Also

Functions
`Simulink.sdi.clearPreferences` | `Simulink.sdi.setRunNamingRule` | `Simulink.sdi.setTableGrouping` | `Simulink.sdi.enablePCTSupport` | `Simulink.sdi.setArchiveRunLimit` | `Simulink.sdi.setAutoArchiveMode`

More About

*   "Iterate Model Design Using the Simulation Data Inspector" (Simulink)
*   "How the Simulation Data Inspector Compares Data" (Simulink)
*   "Compare Simulation Data" (Simulink)
*   "Create Plots Using the Simulation Data Inspector" (Simulink)
*   "Modify Signal Properties in the Simulation Data Inspector" (Simulink)

How the Simulation Data Inspector Compares Data

You can tailor the Simulation Data Inspector comparison process to fit your requirements in multiple ways. When comparing runs, the Simulation Data Inspector:

1   Aligns signal pairs in the **Baseline** and **Compare To** runs based on the **Alignment** settings.

    The Simulation Data Inspector does not compare signals that it cannot align.

2   Synchronizes aligned signal pairs according to the specified **Sync Method**.

    Values for time points added in synchronization are interpolated according to the specified **Interpolation Method**.

3   Computes the difference of the signal pairs.

4   Compares the difference result against specified tolerances.

When the comparison run completes, the results of the comparison are displayed in the navigation pane.

Status	Comparison Result
✅	Difference falls within the specified tolerance.
❌	Difference violates specified tolerance.
➖	The signal does not align with a signal from the **Compare To** run.

When you compare signals with differing time intervals, the Simulation Data Inspector compares the signals on their overlapping interval.

Signal Alignment

In the alignment step, the Simulation Data Inspector decides which signal from the **Compare To** run pairs with a given signal in the **Baseline** run. When you compare signals with the Simulation Data Inspector, you complete the alignment step by selecting the **Baseline** and **Compare To** signals.

The Simulation Data Inspector aligns signals using a combination of their Data Source, Path, SID, and Signal Name properties.

Property	Description
Data Source	Path of the variable in the MATLAB workspace for data imported from the workspace
Path	Block path for the source of the data in its model
SID	Automatically assigned Simulink identifier
Signal Name	Name of the signal in the model

With the default alignment settings, the Simulation Data Inspector aligns signals between runs according to this flow chart.
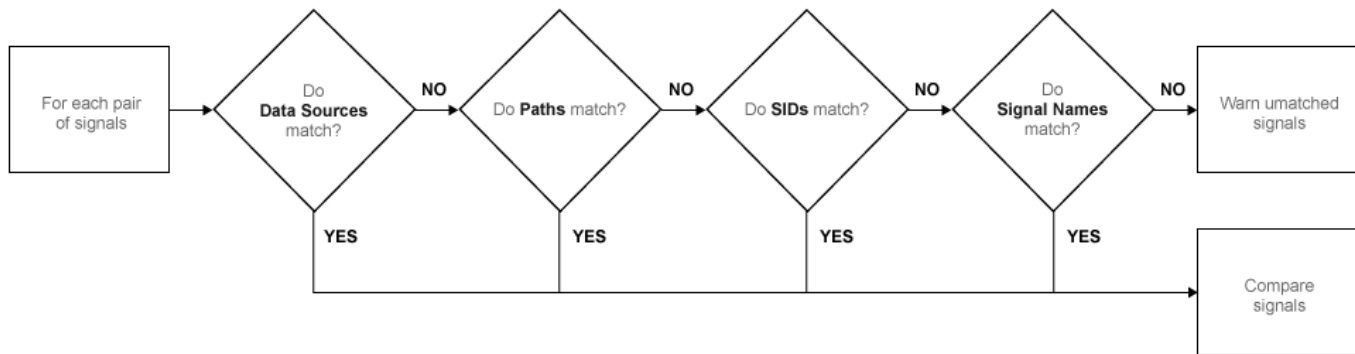
You can specify the priority for each of the signal properties used for alignment in the Simulation Data Inspector **Preferences**. The **Align By** field specifies the highest priority property used to align signals. The priority drops with each subsequent **Then By** field. You must specify a primary alignment property in the **Align By** field, but you can leave any number of the **Then By** fields blank.
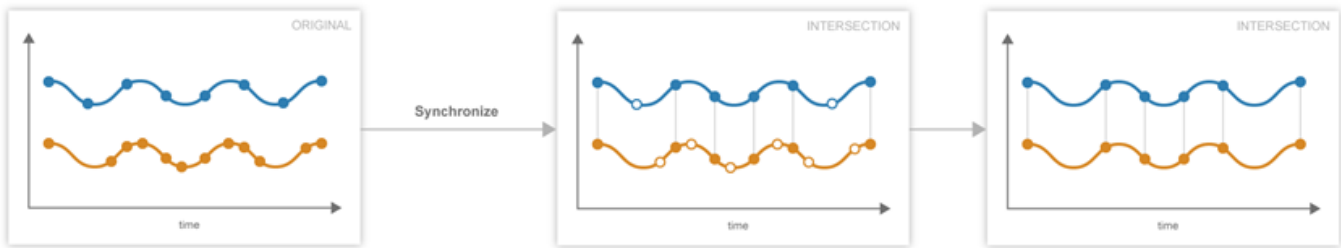
Synchronization

Often, signals that you want to compare don't contain the exact same set of time points. The synchronization step in Simulation Data Inspector comparisons resolves discrepancies in signals' time vectors. You can choose `union` or `intersection` as the synchronization method.

When you specify `union` synchronization, the Simulation Data Inspector builds a time vector that includes every sample time between the two signals. For each sample time not originally present in either signal, the Simulation Data Inspector interpolates the value. The second graph in the illustration shows the union synchronization process, where the Simulation Data Inspector identifies samples to add in each signal, represented by the unfilled circles. The final plot shows the signals after the Simulation Data Inspector has interpolated values for the added time points. The Simulation Data Inspector computes the difference using the signals in the final graph, so that the computed difference signal contains all the data points between the signals.
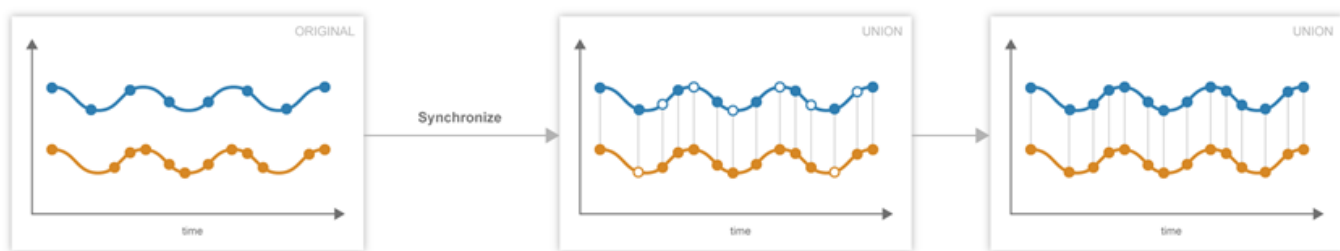
When you specify `intersection` synchronization, the Simulation Data Inspector uses only the sample times present in both signals in the comparison. In the second graph, the Simulation Data Inspector identifies samples that do not have a corresponding sample for comparison, shown as unfilled circles. The final graph shows the signals used for the comparison, without the samples identified in the second graph.

The choice between the synchronization options involves a trade off between speed and accuracy. The interpolation required by `union` synchronization takes time, but provides a more precise result. When you use `intersection` synchronization, the comparison finishes quickly because the Simulation Data Inspector computes the difference for fewer data points and does not interpolate. However, some data is discarded and precision lost with `intersection` synchronization.

Interpolation

The interpolation property of a signal determines how the Simulation Data Inspector displays the signal and how additional data values are computed in synchronization. You can choose to interpolate your data with a zero-order hold (zoh) or a linear approximation. You can also specify no interpolation.
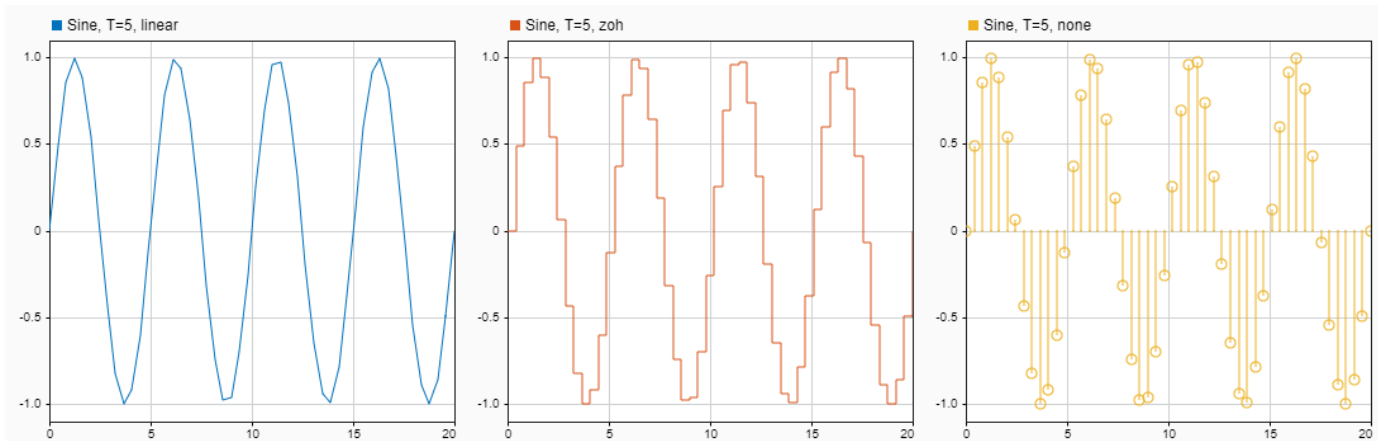
When you specify `zoh` or `none` for the **Interpolation Method**, the Simulation Data Inspector replicates the data of the previous sample for interpolated sample times. When you specify `linear` interpolation, the Simulation Data Inspector uses samples on either side of the interpolated point to linearly approximate the interpolated value. Typically, discrete signals use `zoh` interpolation and continuous signals use `linear` interpolation. You can specify the **Interpolation Method** for your signals in the signal properties.

Tolerance Specification

The Simulation Data Inspector allows you to specify the scope and value of the tolerance for your signal. You can define a tolerance band using any combination of absolute, relative, and time tolerance values, and you can specify whether the specified tolerance applies to an individual signal or to all the signals in a run.

Tolerance Scope

In the Simulation Data Inspector, you can specify the tolerance for your data globally or for an individual signal. Global tolerance values apply to all signals in a run that do not have **Override Global Tol** set to yes. You can specify global tolerance values for your data at the top of the graphical viewing area in the **Compare** view. To specify signal specific tolerance values, edit the signal properties and ensure the **Override Global Tol** property is set to yes.

Tolerance Computation

In the Simulation Data Inspector, you can specify a tolerance band for your run or signal using a combination of absolute, relative, and time tolerance values. When you specify the tolerance for your run or signal using multiple types of tolerances, each tolerance can yield a different answer for the tolerance at each point. The Simulation Data Inspector computes the overall tolerance band by selecting the most lenient tolerance result for each data point.

When you define your tolerance using only the absolute and relative tolerance properties, the Simulation Data Inspector computes the tolerance for each point as a simple maximum.

```
tolerance = max(absoluteTolerance,relativeTolerance*abs(baselineData));
```

The upper boundary of the tolerance band is formed by adding `tolerance` to the **Baseline** signal. Similarly, the Simulation Data Inspector computes the lower boundary of the tolerance band by subtracting `tolerance` from the **Baseline** signal.

When you specify a time tolerance, the Simulation Data Inspector evaluates the time tolerance first, over a time interval defined as $[(t_{samp}\text{-tol}), (t_{samp}\text{+tol})]$ for each sample. The Simulation Data Inspector builds the lower tolerance band by selecting the minimum point on the interval for each sample. Similarly, the maximum point on the interval defines the upper tolerance for each sample.
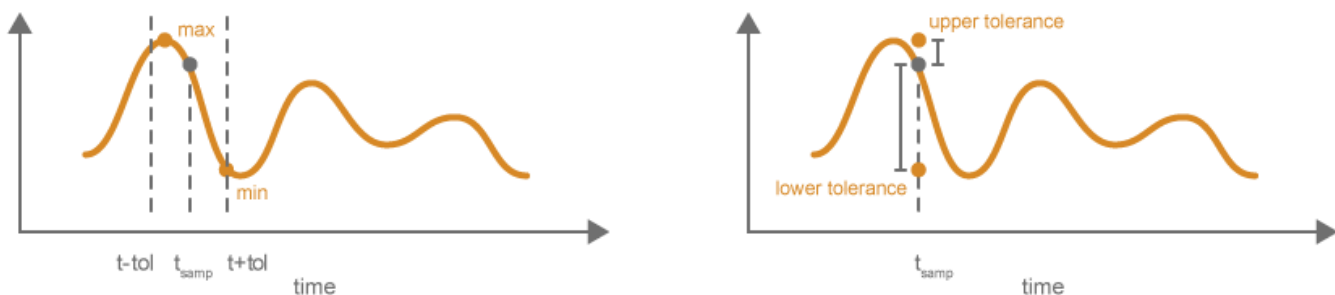
If you specify a tolerance band using an absolute or relative tolerance in addition to a time tolerance, the Simulation Data Inspector applies the time tolerance first, and then applies the absolute and relative tolerances to the maximum and minimum points selected with the time tolerance.
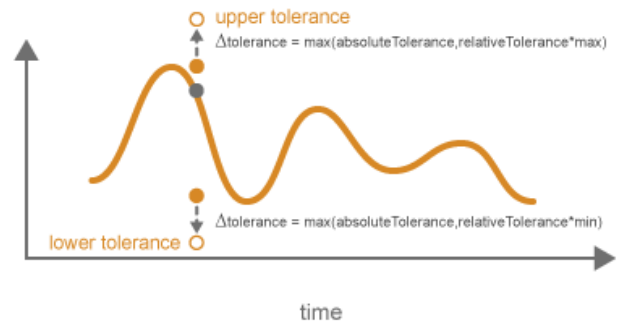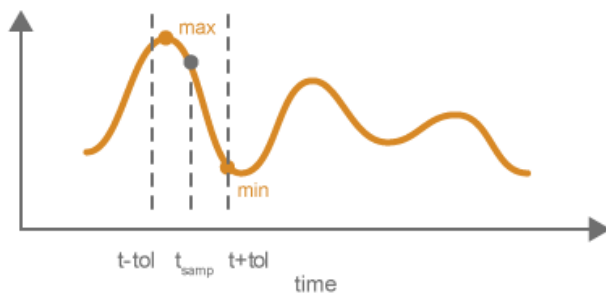
upperTolerance = max + max(absoluteTolerance,relativeTolerance*max)

lowerTolerance = min - max(absoluteTolerance,relativeTolerance*min)

Limitations

The Simulation Data Inspector does not support comparing:

*   Signals of data types `int64` or `uint64`.
*   Variable-size signals.

See Also

Related Examples

*   "Compare Simulation Data" (Simulink)

Save and Share Simulation Data Inspector Data and Views

After you inspect, analyze, or compare your data in the Simulation Data Inspector, you can share your results with others. The Simulation Data Inspector provides several options for sharing and saving your data and results, depending on your needs. With the Simulation Data Inspector, you can:

- Save your data and layout modifications in a Simulation Data Inspector session.
- Share your layout modifications in a Simulation Data Inspector view.
- Share images and figures of plots you create in the Simulation Data Inspector.
- Create a Simulation Data Inspector report.
- Export data to the workspace.
- Export data to a file.

Save and Load Simulation Data Inspector Sessions

If you want to save or share data along with a configured view in the Simulation Data Inspector, save your data and settings in a Simulation Data Inspector session. You can save sessions as MAT- or MLDATX-files. The default format is MLDATX. When you save a Simulation Data Inspector session, the session file contains:

- All runs, data, and properties from the **Inspect** pane, including which run is the current run and which runs are in the archive.
- Plot display selection for signals in the **Inspect** pane.
- Subplot layout and line style and color selections.

Note Comparison results and global tolerances are not saved in Simulation Data Inspector sessions.

To save a Simulation Data Inspector session:

1 Hover over the save icon on the left side bar. Then, click **Save As**.

2 Name the file.
3 Browse to the location where you want to save the session, and click **Save**.

For large datasets, a status overlay in the bottom right of the graphical viewing area displays information about the progress of the save operation and allows you to cancel the save operation.

The **Save** tab of the Simulation Data Inspector preferences menu on the left side bar allows you to configure options related to save operations for MLDATX-files. You can set a limit as low as 50MB on the amount of memory used for the save operation. You can also select one of three **Compression** options:

- None, the default, applies no compression during the save operation.

- `Normal` creates the smallest file size.
- `Fastest` creates a smaller file size than you would get by selecting `None`, but provides a faster save time than `Normal`.

To load a Simulation Data Inspector session, click the open icon on the left side bar. Then, browse to select the MLDATX-file you want to open, and click **Open**.

Alternatively, you can double-click the MLDATX-file. MATLAB and the Simulation Data Inspector open if they are not already open.

When the Simulation Data Inspector already contains runs and you open a session, all of the runs in the session move to the archive. The view updates to show plotted signals from the session file. You can drag runs between the work area and archive as desired.

When the Simulation Data Inspector does not contain runs and you open a session, the Simulation Data Inspector puts runs in the work area and archive as specified in the file.

Share Simulation Data Inspector Views

When you have different sets of data that you want to visualize the same way, you can save a view. A view saves the layout and appearance characteristics of the Simulation Data Inspector without saving the data. Specifically, a view saves:

- Plot visualization type, layout, axis ranges, linking characteristics, and normalized axes
- Location of signals in the plots, including plotted signals in the archive
- Signal grouping and columns on display in the **Inspect** pane
- Signal color and line styling

To save a view:

1
    Click Visualizations and layouts .

2   In **Saved Views**, click **Save current view**.

3   In the dialog box, specify a name for the view and browse to the location where you want to save the MLDATX-file.

4   Click **Save**.

To load a view:

1
    Click Visualizations and layouts .

2   In **Saved Views**, click **Open saved view**.

3   Browse to the view you would like to load, and click **Open**.

Share Simulation Data Inspector Plots

Use the snapshot feature to share the plots you generate in the Simulation Data Inspector. You can export your plots to the clipboard to paste into a document, as an image file, or to a MATLAB figure.

You can choose to capture the entire plot area, including all subplots in the plot area, or to capture only the selected subplot.

Click the camera icon on the toolbar to access the snapshot menu. Use the radio buttons to select the area you want to share and how you want to share the plot. After you make your selections, click **Snapshot** to export the plot.
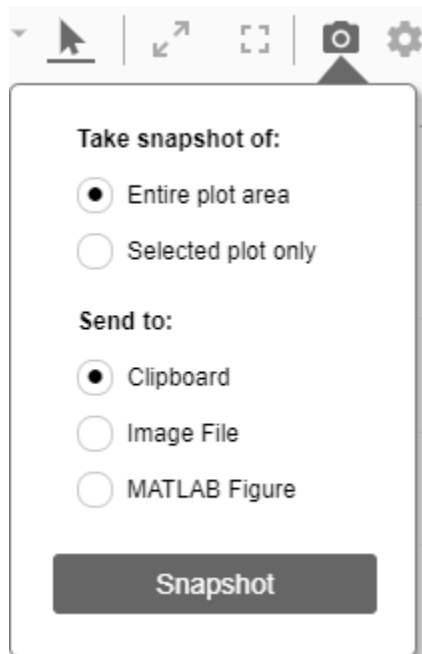
If you create an image, select where you would like to save the image in the file browser.

You can create snapshots of your plots in the Simulation Data Inspector programmatically using `Simulink.sdi.snapshot`.

Create Simulation Data Inspector Report

To generate documentation of your results quickly, create a Simulation Data Inspector report. You can create a report of your data in either the **Inspect** or the **Compare** pane. The report is an HTML file that includes information about all the signals and plots in the active pane. The report includes all signal information displayed in the signal table in the navigation pane. For more information about configuring the table, see "Inspect Metadata" (Simulink).

To generate a Simulation Data Inspector Report:

1

Click the create report icon on the left bar.

2   Specify the type of report you want to create.

 • Select **Inspect** to include the plots and signals from the **Inspect** pane.

- Select **Compare** to include the data and plots from the **Compare** pane. When you generate a **Compare Runs** report, you can choose to **Report only mismatched signals** or to **Report all signals**. If you select **Report only mismatched signals**, the report shows only signal comparisons that are not within the specified tolerances.
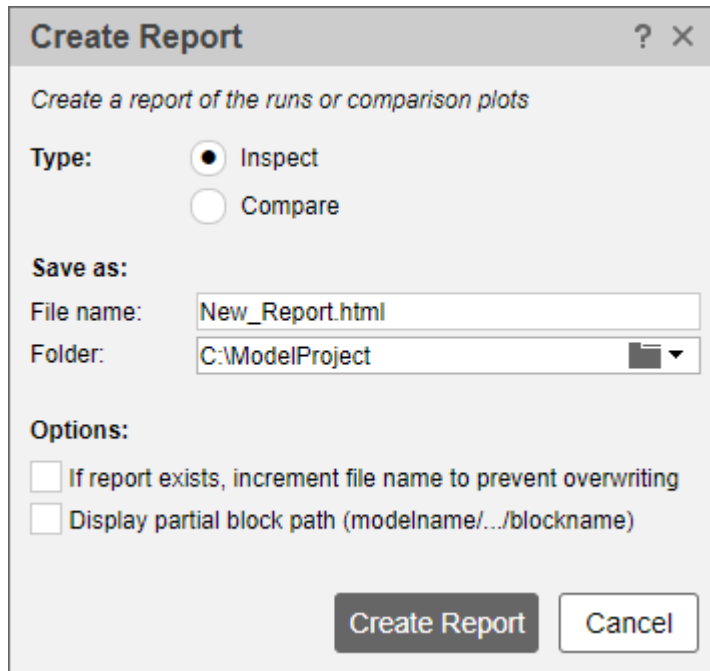
3   Specify a **File name** for the report, and navigate to the **Folder** where you want to save the report.

4   Click **Create Report**.

The generated report automatically opens in your default browser.

Export Data to the Workspace or a File

You can use the Simulation Data Inspector to export data to the base workspace, a MAT file, or a Microsoft Excel file. You can export a selection of runs and signals, runs in the work area, or all runs in the **Inspect** pane, including the **Archive**.

When you export a selection of runs and signals, make the selection of data to export before clicking

the export button .

Only the selected runs and signals are exported. In this example, only the x1 signals from Run 1 and Run 2 are exported. The check box selections for plotting data do not affect whether a signal is exported.
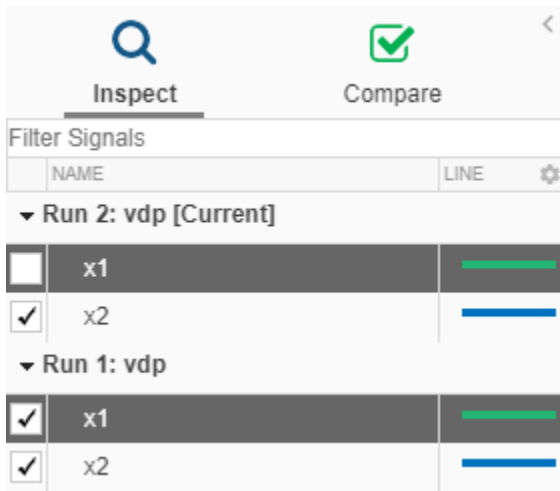
When you export a single signal to the workspace or a MAT file, the signal is exported to a `timeseries` object. Data exported to the workspace or a MAT file for a run or multiple signals is stored as a `Simulink.SimulationData.Dataset` object.

To export data to a file, select the **File** option in the **Export** dialog. You can specify a file name and browse to the location where you want to save the exported file. When you export data to a MAT file, a single exported signal is stored as a `timeseries` object, and runs or multiple signals are stored as a `Simulink.SimulationData.Dataset` object. When you export data to a Microsoft Excel file, the data is stored using the format described in "Microsoft Excel Import, Export, and Logging Format" (Simulink).

To export to a Microsoft Excel file, select the XLSX extension from the drop-down. When you export data to a Microsoft Excel file, you can specify additional options for the format of the data in the exported file. If the file name you provided already exists, you can choose to overwrite the entire file or to only overwrite sheets containing data that corresponds to the exported data. You can also choose which metadata to include and whether signals with identical time data share a time column in the exported file.

Export Video Signal to an MP4 File

You can export a 2D or 3D signal that contains RGB or monochrome video data to an MP4 file using the Simulation Data Inspector. For example, when you log a video signal in a simulation, you can export the data to an MP4 file and view the video using a video player. To export a video signal to an MP4 file:

1   Select the signal you want to export.
2
    Click Export in the toolbar on the left or right-click the signal and select **Export**.
3   In the Export dialog box, choose to export **Selected runs and signals** to a file.
4   Specify a file name and the path to the location where you want to save the file.
5   Select `MP4 video file` from the list and click **Export**.

For the option to export to an MP4 file to be available:

- You must export only one signal at a time.
- The selected signal must be 2D or 3D and contain RGB or monochrome video data.
- The selected signal must be represented in the Simulation Data Inspector as a single signal with multidimensional sample values.

  You may need to convert the signal representation before exporting the signal data. For more information, see "Analyze Multidimensional Signal Data" (Simulink).
- The data type for the signal values must be `double`, `single`, or `uint8`.

Exporting a video signal to an MP4 file is not supported for Linux operating systems.

See Also

Functions
`Simulink.sdi.saveView`

Related Examples

- "View Data in the Simulation Data Inspector" (Simulink)
- "Inspect Simulation Data" (Simulink)
- "Compare Simulation Data" (Simulink)

Inspect and Compare Data Programmatically

You can harness the capabilities of the Simulation Data Inspector from the MATLAB command line using the Simulation Data Inspector API.

The Simulation Data Inspector organizes data in runs and signals, assigning a unique numeric identification to each run and signal. Some Simulation Data Inspector API functions use the run and signal IDs to reference data, rather than accepting the run or signal itself as an input. To access the run IDs in the workspace, you can use `Simulink.sdi.getAllRunIDs` or `Simulink.sdi.getRunIDByIndex`. You can access signal IDs through a `Simulink.sdi.Run` object using the `getSignalIDByIndex` method.

The `Simulink.sdi.Run` and `Simulink.sdi.Signal` classes provide access to your data and allow you to view and modify run and signal metadata. You can modify the Simulation Data Inspector preferences using functions like `Simulink.sdi.setSubPlotLayout`, `Simulink.sdi.setRunNamingRule`, and `Simulink.sdi.setMarkersOn`. To restore the Simulation Data Inspector's default settings, use `Simulink.sdi.clearPreferences`.

Create a Run and View the Data

This example shows how to create a run, add data to it, and then view the data in the Simulation Data Inspector.

Create Data for the Run

Create `timeseries` objects to contain data for a sine signal and a cosine signal. Give each `timeseries` object a descriptive name.

```
time = linspace(0,20,100);

sine_vals = sin(2*pi/5*time);
sine_ts = timeseries(sine_vals,time);
sine_ts.Name = 'Sine, T=5';

cos_vals = cos(2*pi/8*time);
cos_ts = timeseries(cos_vals,time);
cos_ts.Name = 'Cosine, T=8';
```

Create a Run and Add Data

Use the `Simulink.sdi.view` function to open the Simulation Data Inspector.

```
Simulink.sdi.view
```

To import data into the Simulation Data Inspector from the workspace, create a `Simulink.sdi.Run` object using the `Simulink.sdi.Run.create` function. Add information about the run to its metadata using the `Name` and `Description` properties of the `Run` object.

```
sinusoidsRun = Simulink.sdi.Run.create;
sinusoidsRun.Name = 'Sinusoids';
sinusoidsRun.Description = 'Sine and cosine signals with different frequencies';
```

Use the `add` function to add the data you created in the workspace to the empty run.

```
add(sinusoidsRun,'vars',sine_ts,cos_ts);
```

Plot the Data in the Simulation Data Inspector

Use the `getSignalByIndex` function to access `Simulink.sdi.Signal` objects that contain the signal data. You can use the `Simulink.sdi.Signal` object properties to specify the line style and color for the signal and plot it in the Simulation Data Inspector. Specify the `LineColor` and `LineDashed` properties for each signal.

```
sine_sig = getSignalByIndex(sinusoidsRun,1);
sine_sig.LineColor = [0 0 1];
sine_sig.LineDashed = '-.';

cos_sig = sinusoidsRun.getSignalByIndex(2);
cos_sig.LineColor = [0 1 0];
cos_sig.LineDashed = '--';
```

Use the `Simulink.sdi.setSubPlotLayout` function to configure a 2-by-1 subplot layout in the Simulation Data Inspector plotting area. Then use the `plotOnSubplot` function to plot the sine signal on the top subplot and the cosine signal on the lower subplot.

```
Simulink.sdi.setSubPlotLayout(2,1);

plotOnSubPlot(sine_sig,1,1,true);
plotOnSubPlot(cos_sig,2,1,true);
```

Close the Simulation Data Inspector and Save Your Data

When you have finished inspecting the plotted signal data, you can close the Simulation Data Inspector and save the session to an MLDATX file.

```
Simulink.sdi.close('sinusoids.mldatx')
```

Compare Two Signals in the Same Run

You can use the Simulation Data Inspector programmatic interface to compare signals within a single run. This example compares the input and output signals of an aircraft longitudinal controller.

First, load the session that contains the data.

```
Simulink.sdi.load('AircraftExample.mldatx');
```

Use the `Simulink.sdi.Run.getLatest` function to access the latest run in the data.

```
aircraftRun = Simulink.sdi.Run.getLatest;
```

Then, you can use the `Simulink.sdi.getSignalsByName` function to access the `Stick` signal, which represents the input to the controller, and the `alpha, rad` signal that represents the output.

```
stick = getSignalsByName(aircraftRun,'Stick');
alpha = getSignalsByName(aircraftRun,'alpha, rad');
```

Before you compare the signals, you can specify a tolerance value to use for the comparison. Comparisons use tolerance values specified for the baseline signal in the comparison, so set an absolute tolerance value of `0.1` on the `Stick` signal.

```
stick.AbsTol = 0.1;
```

Now, compare the signals using the `Simulink.sdi.compareSignals` function. The `Stick` signal is the baseline, and the `alpha, rad` signal is the signal to compare against the baseline.

```
comparisonResults = Simulink.sdi.compareSignals(stick.ID,alpha.ID);
match = comparisonResults.Status

match =
  ComparisonSignalStatus enumeration

    OutOfTolerance
```

The comparison result is out of tolerance. You can use the `Simulink.sdi.view` function to open the Simulation Data Inspector to view and analyze the comparison results.

Compare Runs with Global Tolerance

You can specify global tolerance values to use when comparing two simulation runs. Global tolerance values are applied to all signals within the run. This example shows how to specify global tolerance values for a run comparison and how to analyze and save the comparison results.

First, load the session file that contains the data to compare. The session file contains data for four simulations of an aircraft longitudinal controller. This example compares data from two runs that use different input filter time constants.

```
Simulink.sdi.load('AircraftExample.mldatx');
```

To access the run data to compare, use the `Simulink.sdi.getAllRunIDs` (Simulink) function to get the run IDs that correspond to the last two simulation runs.

```
runIDs = Simulink.sdi.getAllRunIDs;
runID1 = runIDs(end - 1);
runID2 = runIDs(end);
```

Use the `Simulink.sdi.compareRuns` (Simulink) function to compare the runs. Specify a global relative tolerance value of `0.2` and a global time tolerance value of `0.5`.

```
runResult = Simulink.sdi.compareRuns(runID1,runID2,'reltol',0.2,'timetol',0.5);
```

Check the `Summary` property of the returned `Simulink.sdi.DiffRunResult` object to see whether signals were within the tolerance values or out of tolerance.

```
runResult.Summary

ans = struct with fields:
       OutOfTolerance: 0
      WithinTolerance: 3
            Unaligned: 0
        UnitsMismatch: 0
                Empty: 0
             Canceled: 0
          EmptySynced: 0
     DataTypeMismatch: 0
```

```
        TimeMismatch: 0
   StartStopMismatch: 0
         Unsupported: 0
```

All three signal comparison results fell within the specified global tolerance.

You can save the comparison results to an MLDATX file using the `saveResult` (Simulink) function.

```
saveResult(runResult,'InputFilterComparison');
```

Analyze Simulation Data Using Signal Tolerances

You can programmatically specify signal tolerance values to use in comparisons performed using the Simulation Data Inspector. In this example, you compare data collected by simulating a model of an aircraft longitudinal flight control system. Each simulation uses a different value for the input filter time constant and logs the input and output signals. You analyze the effect of the time constant change by comparing results using the Simulation Data Inspector and signal tolerances.

First, load the session file that contains the simulation data.

```
Simulink.sdi.load('AircraftExample.mldatx');
```

The session file contains four runs. In this example, you compare data from the first two runs in the file. Access the `Simulink.sdi.Run` objects for the first two runs loaded from the file.

```
runIDs = Simulink.sdi.getAllRunIDs;
runIDTs1 = runIDs(end-3);
runIDTs2 = runIDs(end-2);
```

Now, compare the two runs without specifying any tolerances.

```
noTolDiffResult = Simulink.sdi.compareRuns(runIDTs1,runIDTs2);
```

Use the `getResultByIndex` function to access the comparison results for the `q` and `alpha` signals.

```
qResult = getResultByIndex(noTolDiffResult,1);
alphaResult = getResultByIndex(noTolDiffResult,2);
```

Check the `Status` of each signal result to see whether the comparison result fell within our out of tolerance.

```
qResult.Status
```

```
ans =
  ComparisonSignalStatus enumeration

    OutOfTolerance
```

```
alphaResult.Status
```

```
ans =
  ComparisonSignalStatus enumeration
```

```
OutOfTolerance
```

The comparison used a value of `0` for all tolerances, so the `OutOfTolerance` result means the signals are not identical.

You can further analyze the effect of the time constant by specifying tolerance values for the signals. Specify the tolerances by setting the properties for the `Simulink.sdi.Signal` objects that correspond to the signals being compared. Comparisons use tolerances specified for the baseline signals. This example specifies a time tolerance and an absolute tolerance.

To specify a tolerance, first access the `Signal` objects from the baseline run.

```
runTs1 = Simulink.sdi.getRun(runIDTs1);
qSig = getSignalsByName(runTs1,'q, rad/sec');
alphaSig = getSignalsByName(runTs1,'alpha, rad');
```

Specify an absolute tolerance of `0.1` and a time tolerance of `0.6` for the `q` signal using the `AbsTol` and `TimeTol` properties.

```
qSig.AbsTol = 0.1;
qSig.TimeTol = 0.6;
```

Specify an absolute tolerance of `0.2` and a time tolerance of `0.8` for the `alpha` signal.

```
alphaSig.AbsTol = 0.2;
alphaSig.TimeTol = 0.8;
```

Compare the results again. Access the results from the comparison and check the `Status` property for each signal.

```
tolDiffResult = Simulink.sdi.compareRuns(runIDTs1,runIDTs2);
qResult2 = getResultByIndex(tolDiffResult,1);
alphaResult2 = getResultByIndex(tolDiffResult,2);
```

`qResult2.Status`

```
ans =
  ComparisonSignalStatus enumeration

    WithinTolerance
```

`alphaResult2.Status`

```
ans =
  ComparisonSignalStatus enumeration

    WithinTolerance
```

See Also
Simulation Data Inspector

Related Examples

- "Compare Simulation Data" (Simulink)
- "How the Simulation Data Inspector Compares Data" (Simulink)
- "Create Plots Using the Simulation Data Inspector" (Simulink)

Limit the Size of Logged Data

In this section...
"Limit the Number of Runs Retained in the Simulation Data Inspector Archive" on page 5-48
"Specify a Minimum Disk Space Requirement or Maximum Size for Logged Data" on page 5-48
"View Data Only During Simulation" on page 5-49
"Reduce the Number of Data Points Logged from Simulation" on page 5-49

Logging simulation data can produce large amounts of data that fill up disk space. Such situations include logging many signals, logging data for long simulations, and running many simulations without deleting run data from the Simulation Data Inspector. You can choose among several options to limit the size of logged simulation data. You can:

- Limit the number of runs retained in the Simulation Data Inspector archive.
- Reduce the number of data points logged in each simulation.
- Specify a minimum disk space requirement or maximum size for logged data.
- Configure logging for only viewing data during simulation.

Depending on your requirements, you can use more than one strategy to limit the size of logged data.

Limit the Number of Runs Retained in the Simulation Data Inspector Archive

When you run multiple simulations in a single MATLAB session, logged simulation data accumulates in the Simulation Data Inspector even if you overwrite the logging data in the MATLAB workspace. To reduce the amount of data the Simulation Data Inspector retains, you can configure a limit for the number of runs stored in the archive. When the number of runs in the archive reaches the size limit, the Simulation Data Inspector starts to delete runs from the archive on a first-in, first-out basis.

Configure the archive **Size** setting in the Simulation Data Inspector preferences. The size limit only applies to runs in the archive. For the Simulation Data Inspector to automatically limit data retention, select **Automatically archive** and specify the maximum number of runs to retain in the archive. By default, **Automatically archive** is enabled with an archive size limit of twenty runs. If you experience issues with logged data consuming too much disk space, consider adjusting the size limit for the archive in the Simulation Data Inspector preferences.

Specify a Minimum Disk Space Requirement or Maximum Size for Logged Data

You can use preferences in the Simulation Data Inspector to directly limit the size of logged data by specifying a minimum amount of disk space to leave free or by specifying a maximum size for logged data on disk. Each setting accounts for all kinds of logged data. By default, logged data must leave at least 100 MB of free disk space with no maximum size limit. Specify the required disk space and maximum size in GB, and specify 0 to apply no disk space requirement or no maximum size limit.

When you specify a minimum disk space requirement or a maximum size for logged data, you can also specify whether to prioritize retaining data from the current simulation or data from prior simulations when approaching the limit. By default, the Simulation Data Inspector prioritizes

retaining data for the current run. As the free disk space or logged data size approaches the limit, prior runs are deleted first to free up space for data being logged in the current run. If deleting runs does not free up enough space, recording is disabled. To prioritize retaining prior data, change the **When low on disk space** setting to **Keep prior runs and stop recording**. You see a warning message when prior runs are deleted and when recording is disabled. If recording is disabled due to the size of logged data, you need to change the **Record Mode** back to **View and record data** to continue logging data, after you have freed up disk space.

View Data Only During Simulation

In some situations, you may want to only view the data for logged signals and not save the values. For example, when using the Simulation Data Inspector to visualize data streaming from hardware, you may only want to view the data live and not record it. You can change the **Record mode** in the Simulation Data Inspector preferences to **View during simulation only** so that logged data is not saved and you can still view the data during simulation. The **Record mode** is reset to **View and record data** at the start of each MATLAB session.

When you change the **Record mode** to **View during simulation only**:

- Logged data is not available in the Simulation Data Inspector or workspace after simulation.
- You can view data using dashboard blocks, scopes, and the Simulation Data Inspector, but plots clear when you pan or zoom.
- You cannot access logged data during simulation using the Simulation Data Inspector programmatic interface.

Reduce the Number of Data Points Logged from Simulation

Model configuration parameters and signal properties allow you to limit the number of data points logged in a simulation. Be sure to carefully consider data requirements when limiting logged data points. Limiting data can skip critical time points, and can lead to aliasing, if your effective sample rate is too low.

You can reduce the number of data points using:

- Decimation — Log every nth signal value.
- Limit data points to last — Only log the last n signal values.
- Logging intervals — Specify specific time intervals in which to log data.

For details, see "Specify Signal Values to Log" (Simulink).

See Also

Tools
Simulation Data Inspector

Related Examples

- "Specify Signal Values to Log" (Simulink)
- "Configure the Simulation Data Inspector" (Simulink)